LSI

Storage. Networking. Accelerated.™

Syncro™ CS 9286-8e Solution

User Guide

Version 1.0
March 2013

DB15-001017-00

Revision History

Version and Date	Description of Changes
Version 1.0, March 2013	Initial release of this document.

Corporate Headquarters
San Jose, CA
800-372-2447

Email
globalsupport@lsi.com

Website
www.lsi.com

Document Number: DB15-001017-00

Table of Contents

Chapter 1: Introduction

This document explains how to set up and configure the hardware and software for the Syncro™ CS 9286-8e high-availability direct-attached storage (HA-DAS) solution.

The Syncro CS 9286-8e solution provides fault tolerance capabilities as a key part of a high-availability data storage system. The Syncro CS 9286-8e solution combines redundant servers, LSI® HA-DAS RAID controllers, computer nodes, cable connections, common SAS JBOD enclosures, and dual-ported SAS storage devices.

The redundant components and software technologies provide a high-availability system with ongoing service that is not interrupted by the following events:

- The failure of a single internal node does not interrupt service because the solution has multiple nodes with cluster failover.
- An expander failure does not interrupt service because the dual expanders in every enclosure provide redundant data paths.
- A drive failure does not interrupt service because RAID fault tolerance is part of the configuration.
- A system storage expansion or maintenance activity can be completed without requiring an interruption of service because of redundant components, management software, and maintenance procedures.

1.1 Concepts of High-Availability DAS

In terms of data storage and processing, *High Availability* (HA) means a computer system design that ensures a high level of operational continuity and data access reliability over a long period of time. High-availability systems are critical to the success and business needs of small and medium-sized business (SMB) customers, such as retail outlets and health care offices, who cannot afford to have their computer systems go down. An HA-DAS solution enables customers to maintain continuous access to and use of their computer system. Shared direct-attached drives are accessible to multiple servers, thereby maintaining ease of use and reducing storage costs.

A *cluster* is a group of computers working together to run a common set of applications and to present a single logical system to the client and application. *Failover clustering* provides redundancy to the cluster group to maximize up-time by utilizing fault-tolerant components. In the example of two servers with shared storage that comprise a failover cluster, when a server fails, the failover cluster automatically moves control of the shared resources to the surviving server with no interruption of processing. This configuration allows seamless failover capabilities in the event of planned failover (maintenance mode) for maintenance or upgrade, or in the event of a failure of the CPU, memory, or other server failures.

Because multiple initiators exist in a clustered pair of servers (nodes) with a common shared storage domain, there is a concept of *device reservations* in which physical drives, drive groups, and virtual drives (VDs) are managed by a selected single initiator. For HA-DAS, I/O transactions and RAID management operations are normally processed by a single Syncro CS 9286-8e controller, and the associated physical drives, drive groups, and VDs are only visible to that controller. To assure continued operation, all other physical drives, drive groups, and VDs are also visible to, though not normally controlled by, the Syncro CS controller. This key functionality allows the Syncro CS 9286-8e solution to share VDs among multiple initiators as well as exclusively constrain VD access to a particular initiator without the need for SAS zoning.

Node downtime in an HA system can be either *planned* and *unplanned*. *Planned* node downtime is the result of management-initiated events, such as upgrades and maintenance. *Unplanned* node downtime results from events that are not within the direct control of IT administrators, such as failed software, drivers, or hardware. The Syncro CS 9286-8e solution protects your data and maintains system up-time from both planned and unplanned node downtime. It also enables you to schedule node downtime to update hardware or firmware, and so on. When you bring one controller node down for scheduled maintenance, the other node takes over with no interruption of service.

1.2 HA-DAS Terminology

This section defines some additional important HA-DAS terms.

- **Cache Mirror**: A cache coherency term describing the duplication of write-back cached data across two controllers.
- **Exclusive Access**: A host access policy in which a VD is only exposed to, and accessed by, a single specified server.
- **Failover**: The process in which the management of drive groups and VDs transitions from one controller to the peer controller to maintain data access and availability.
- **HA Domain**: A type of storage domain that consists of a set of HA controllers, cables, shared disk resources, and storage media.
- **Peer Controller**: A relative term to describe the HA controller in the HA domain that acts as the failover controller.
- **Server/Controller Node**: A processing entity composed of a single host processor unit or multiple host processor units that is characterized by having a single instance of a host operating system.
- **Server Storage Cluster**: An HA storage topology in which a common pool of storage devices is shared by two computer nodes through dedicated Syncro CS 9286-8e controllers.
- **Shared Access**: A host access policy in which a VD is exposed to, and can be accessed by, all servers in the HA domain.
- **Virtual Drive (VD)**: A storage unit created by a RAID controller from one or more physical drives. Although a virtual drive can consist of multiple drives, it is seen by the operating system as a single drive. Depending on the RAID level used, the virtual drive might retain redundant data in case of a drive failure.

1.3 Syncro CS 9286-8e Solution Features

The Syncro CS 9286-8e solution supports the following HA features.

- Server storage cluster topology, enabled by the following supported operating systems:
  - Microsoft® Windows® Server 2008 R2
  - Microsoft Windows Server 2012
- Clustering/HA services support:
  - Microsoft failover clustering
- Dual-active HA with shared storage
- Controller-to-controller intercommunication over SAS
- Write-back cache coherency
- CacheCade 1.0 (Read)
- Shared and exclusive VD I/O access policies
- Operating system boot from the controller (exclusive access)
- Controller hardware and property mismatch detection, handling, and reporting
- Global hot spare support for all volumes in the HA domain
- Planned and unplanned failover modes
- CacheVault™ provides cache cached data protection in case of host power loss or server failure
- Full MegaRAID® features, with the following exceptions.
  - T10 Data Integrity Field (DIF) is not supported.
  - Self-encrypting drives (SED) and full disk encryption (FDE) are not supported.
  - CacheCade 2.0 (write back) is not supported.
  - Dimmer switch is not supported.
  - SGPIO sideband signaling for enclosure management is not supported.

1.4 Hardware Compatibility

The servers, disk drives, and JBOD enclosures you use in the Syncro CS 9286-8e solution must be selected from the list of approved components that LSI has tested for compatibility. Refer to this web link for the compatibility lists.

http://www.lsi.com/channel/support/pages/interoperability.aspx

Chapter 2: Hardware and Software Setup

This chapter explains how to set up the hardware and software for a Syncro CS 9286-8e solution with two controller nodes and shared storage. For this implementation you use two standard server modules with Syncro CS 9286-8e controllers that provide access to disks in one or more JBOD enclosures for reliable, high-access redundancy.

The LSI Syncro CS 9286-8e controller kit includes the following items:

- Two Syncro CS 9286-8e controllers
- Two CacheVault Flash Module 03 (CVFM03) devices (pre-installed on the controllers)
- Two CacheVault Power Module 02 (CVPM02) devices
- Two CVPM02 mounting clips and hardware
- Two CVPM02 extension cables
- Two low-profile brackets
- *Syncro CS 9286-8e Controller Quick Installation Guide*
- Syncro CS Resource CD

The hardware configuration for the Syncro CS 9286-8e solution requires the following additional hardware that is not included in the kit:

- Two server modules from the LSI-approved compatibility list. The servers must include network cards.
- One or more JBOD enclosures with SAS disk drives from the LSI-approved compatibility list.
- A monitor and mouse for each server node.
- Network cabling and SAS cabling to connect the servers and JBOD enclosures.

The following figure shows a high-level diagram of a Syncro CS 9286-8e solution connected to a network.

Figure 1 Two-Server Syncro CS 9286-8e Configuration

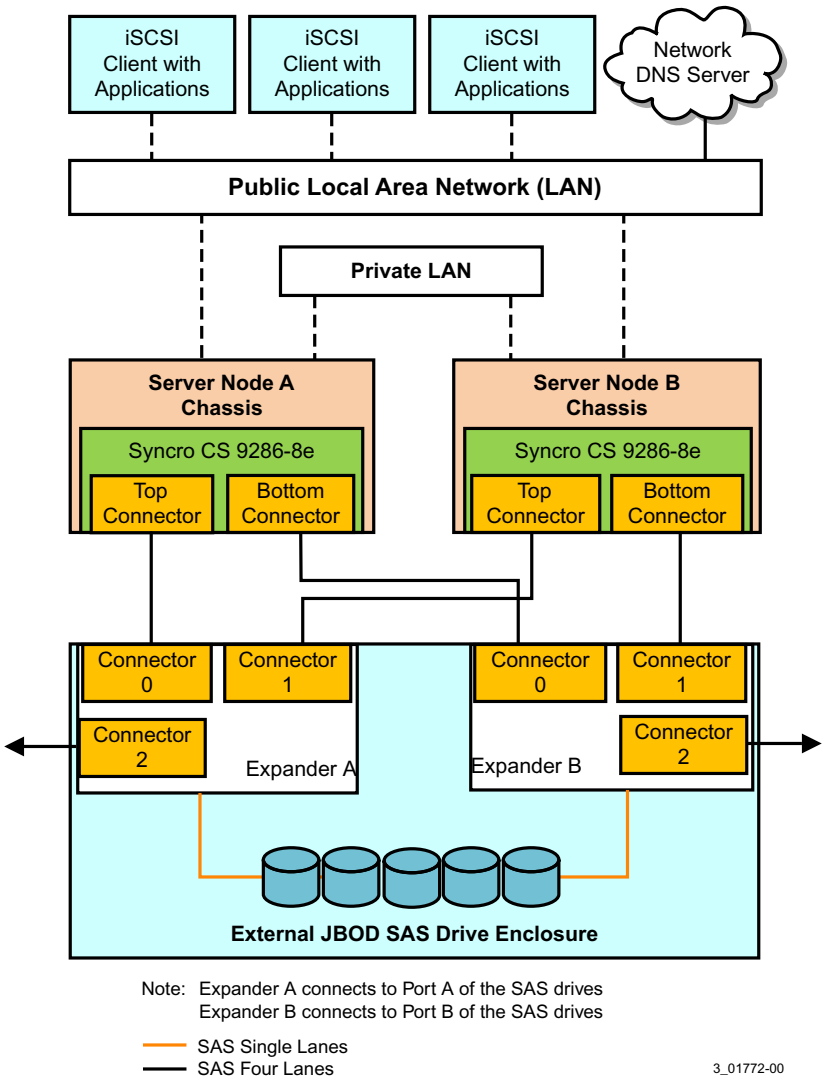

2.1 Setting Up a Syncro CS 9286-8e Two-Server Cluster Configuration

Follow these steps to set up the hardware for a Syncro CS 9286-8e configuration.

1. Unpack the Syncro CS 9286-8e controllers and the CVPM02 modules from the kit and inspect them for damage.

   If any components of the kit appear to be damaged, or if any items are missing from the kit, contact your LSI Customer and Technical Support representative.
2. Turn off the power to the server units, disconnect the power cords, and disconnect any network cables.
3. Remove the covers from the two server units.

   Refer to the server documentation for instructions on how to do this.
4. Review the Syncro CS 9286-8e jumper settings and change them if necessary. Also note the location of the two external Mini-SAS SFF8088 connectors J1A4 and J1B1.

   You usually *do not* need to change the default factory settings of the jumpers. The following figure shows the location of the jumpers and connectors on the controller board.

The CVFM03 module comes preinstalled on the Syncro CS controller; however, the module is not included in the following figure so that you can see all of the connectors and headers on the controller board. Figure 3 and Figure 4 show the controller with the CVFM03 module installed.

Figure 2 Syncro CS 9286-8e Jumpers and Connectors

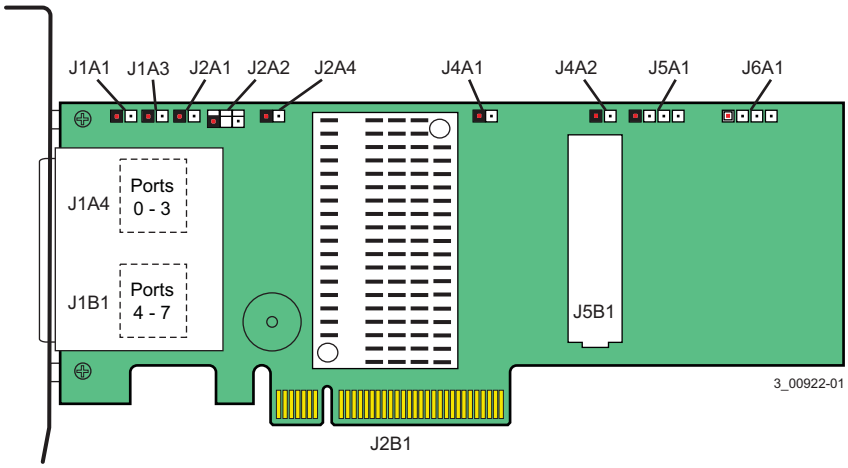

In the figure, Pin 1 is highlighted in red for each jumper. The following table describes the jumpers and connectors on the Syncro CS 9286-8e controller.

Table 1 Syncro CS 9286-8e Controller Jumpers

Jumper/ Connector	Type	Description
J1A4	External SFF-8088 4-port SAS connector	In the cabling figures later in this chapter, this connector is called the "top" connector.
J1B1	External SFF-8088 4-port SAS connector	In the cabling figures later in this chapter, this connector is called the "bottom" connector.
J1A1	Write-Pending LED header	2-pin connector Connects to an LED that indicates when the data in the cache has yet to be written to the storage devices. Used when the write-back feature is enabled.
J1A3	Global Drive Fault LED header	2-pin connector Connects to an LED that indicates whether a drive is in a fault condition.
J2A1	Activity LED header	2-pin connector Connects to an LED that indicates activity on the drives connected to the controller.
J2A2	Advanced Software Options Hardware Key header	3-pin header Enables support for the Advanced Software Options features.
J2A4	I^2O Mode jumper	2-pin connector Installing this jumper causes the RAID controller to run in I^2O mode. The default, recommended mode of operation is without the shunt and running in Fusion mode.
J4A1	Serial EEPROM	2-pin connector Provides controller information, such as the serial number, revision, and manufacturing date. The default is no shunt installed.

Table 1 Syncro CS 9286-8e Controller Jumpers (Continued)

Jumper/ Connector	Type	Description
J4A2	LSI Test header	Reserved for LSI use.
J5A1	Serial UART connector for the expander	Reserved for LSI use.
J6A1	Serial UART connector for the expander	Reserved for LSI use.

5. Place the Syncro CS controller on a flat, clean, static-free surface after grounding yourself.
6. Take the cable included with the kit and insert one end of it into the 9-pin connector on the remote CVPM02 module, as shown in the following figure.

NOTE The CVPM02 module is a super-capacitor pack that provides power for the cache offload capability to protect cached data in case of host power loss or server failure.

Figure 3 Connecting the Cable to the Remote CVPM02 Module

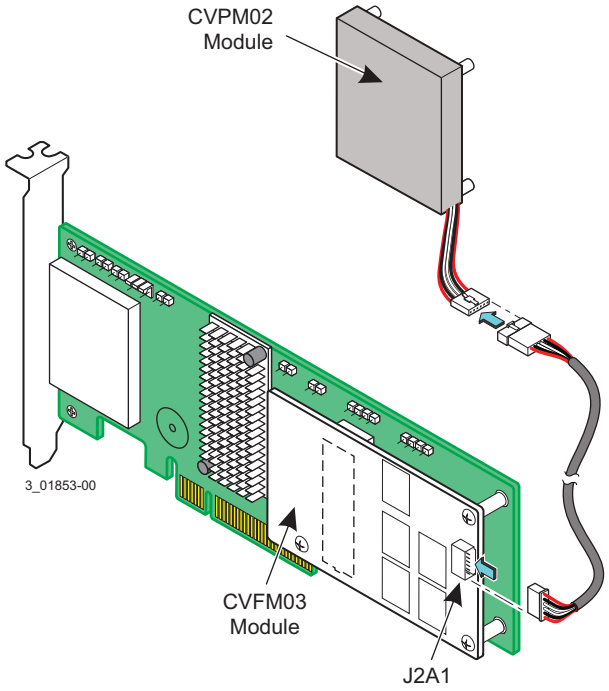

7. Mount the CVPM02 module to the inside of the server chassis, as shown in Figure 4. Refer to the server documentation to determine the exact method of mounting the module.

NOTE Because server chassis design varies from vendor to vendor, no standard mounting option exists for the CVPM02 module that is compatible with all chassis configurations. Authorized resellers and chassis manufacturers might have recommendations for the location of the power module to provide the most flexibility within various environments.

8. Make sure that the power to the server is still turned off, that the power cords are unplugged, and that the chassis is grounded and has no AC power.
9. Insert the Syncro CS 9286-8e controller into a PCIe slot on the motherboard, as shown in the following figure.

Press down gently, but firmly, to seat the Syncro CS 9286-8e controller correctly in the slot.

NOTE The Syncro CS 9286-8e controller is a PCIe *x*8 card that can operate in *x*8 or *x*16 slots. Some x16 PCIe slots support only PCIe graphics cards; if you install a Syncro CS 9286-8e in one of these slots, the controller will not function. Refer to the motherboard documentation for information about the configuration of the PCIe slots.

Figure 4 Installing the Syncro CS 9286-8e Controller and Connecting the Cable

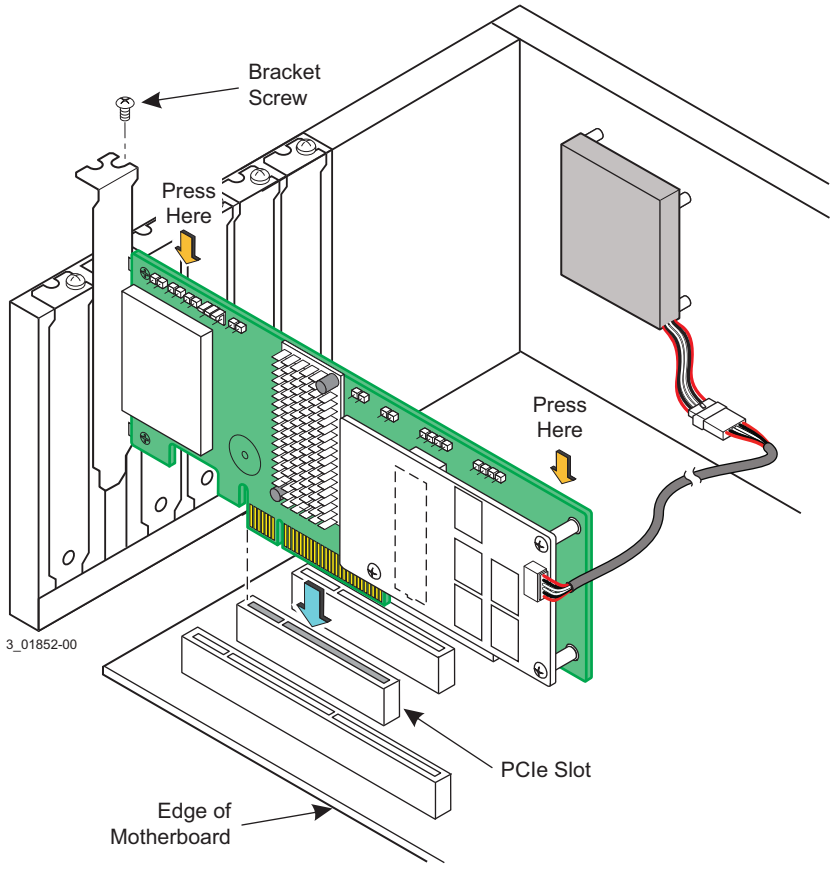

10. Secure the controller to the computer chassis with the bracket screw.
11. Insert the other 9-pin cable connector on the cable into the J2A1 connector on the CVFM03 module, as shown in Figure 3.
12. Repeat step 5 to step 11 to install the second Syncro CS 9286-8e controller in the second server module.
13. Install SAS disk drives in the JBOD enclosure or enclosures.

NOTE In a Syncro CS configuration, the expanders in the JBOD enclosure must have two four-lane IN ports. As an option, the expanders can be configured with a third four-lane port to connect to other cascaded expanders, as shown in Figure 1. JBOD enclosures with dual expanders can support *split* mode or *unified* mode. For fault-tolerant cabling configurations, you typically configure the JBOD enclosure in unified mode. (Check with the JBOD enclosure vendor to determine the appropriate settings).

Refer to the drive documentation to determine any pre-installation configuration requirements. Be sure to use SAS disk drives that are listed on the LSI-approved list. (To view this list, follow the URL listed in Section 1.4, Hardware Compatibility.)

14. If necessary, install network boards in the two server modules and install the cabling between them.
15. Reinstall the covers of the two servers.
16. Install the two server modules and the JBOD enclosure in an industry-standard cabinet, if appropriate, following the manufacturer's instructions.
17. Use SAS cables to connect the two external connectors on the Syncro CS 9286-8e controller to the JBOD enclosure or enclosures. See Figure 2 to view the location of the external connectors.

    See Section 2.2, Cabling Configurations, for specific cabling instructions for one or two JBOD enclosures.
18. Reconnect the power cords and turn on the power to the servers and the JBOD enclosure or enclosures.

    Follow the generally accepted best practice by turning on the power on to the JBOD enclosure before you power the two servers. If you power the servers before you power the JBOD enclosure, the servers might not recognize the disk drives.

    When the servers boot, a BIOS message appears. The firmware takes several seconds to initialize. The configuration utility prompt times out after several seconds. The second portion of the BIOS message shows the Syncro CS 9286-8e number, firmware version, and cache SDRAM size. The numbering of the controllers follows the PCI slot scanning order used by the host motherboard.
19. Configure the groups and the virtual drives on the two controllers.

    For specific instructions, see Section 3.3, Creating Virtual Drives on the Controller Nodes. You can use WebBIOS, MegaCLI64, or MegaRAID Storage Manager to create the groups and virtual drives.
20. Install the operating system driver on both server nodes.

    You must install the software drivers first, before installing the operating system.

    You can view the supported operating systems and download the latest drivers for the Syncro CS controllers from the LSI website at http://www.lsi.com/support/Pages/download-search.aspx. Access the download center, and follow the steps to download the appropriate driver.

    Refer to the *MegaRAID SAS Device Driver Installation User Guide* on the Syncro CS Resource CD for more information about installing the driver. Be sure to review the readme file that accompanies the driver.
21. Install the operating system on both server nodes, following the instructions from Microsoft. The following operating systems are supported:
    - Microsoft Windows Server 2012
    - Microsoft Windows Server 2008 R2

    NOTE Support for clustered RAID controllers is not enabled by default in Microsoft Windows Server 2012 or Microsoft Windows Server 2008 R2. To enable support for this feature, consult with your *server vendor. For additional information visit the Cluster* in a Box Validation Kit for Windows Server site on the Microsoft Windows Server TechCenter website.

    Be sure to use the latest service packs that are provided by Microsoft. You have two options for installing the operating system for each controller node:
    - Install it on a private volume connected to the system-native storage controller. The recommended best practice is to install the operating system on this private volume because the disks in the clustering configuration cannot see this volume. Therefore, no danger exists of accidentally overwriting the operating system disk when you set up clustering.
    - Install it on an exclusive virtual drive connected to the Syncro CS 9286-8e controller. Exclusive host access is required for a boot volume so the volume is not overwritten accidentally when you create virtual drives for data storage. For instructions on creating exclusive virtual drives using the WebBIOS utility, see Section 3.3.1, Creating Shared or Exclusive VDs with the WebBIOS Utility.

NOTE The Syncro CS 9286-8e solution does not support booting from a shared operating system volume.

22. Install the Failover Cluster feature on both servers, following the instructions in the Microsoft documentation.

2.2 Cabling Configurations

This section has information about initially setting up a Syncro CS configuration with one or two JBOD enclosures. It also explains how to add a second JBOD enclosure to an operational single-JBOD configuration without interrupting service on the configuration.

System throughput problems can occur if you use the wrong kind of SAS cables. To minimize the potential for problems, use high-quality cables that meet SAS 2.1 standards and that are less than 6 meters long. See the list of approved cables and vendors on the web link listed at the end of Section 1.4, Hardware Compatibility.

The following figure shows the SAS cable connections for a two-controller-node configuration with a single JBOD enclosure.

NOTE In the figures in this section, *Top Connector* means the external connector closest to the bracket screw support end of the controller bracket (connector J1A4 in Figure 2). *Bottom Connector* means the other external connector (connector J1B1 in Figure 2).

Figure 5 Two-Controller-Node Configuration with Single JBOD Enclosure

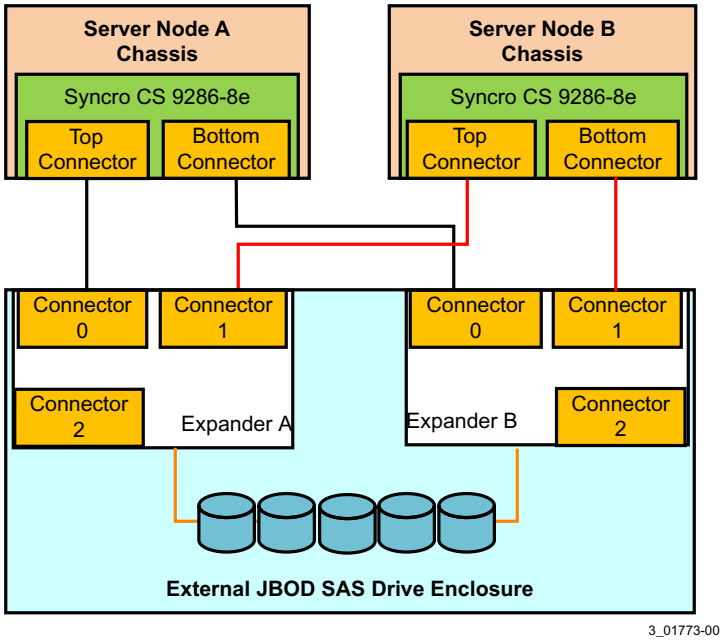

The cross-connections between the controllers provide redundant paths that safeguard against expander, cable, or expander failure.

To retain consistent device reporting, the corresponding port numbers for both controllers must be connected to a common enclosure expander. In this example, the top connector on each controllers is connected to Expander A and the bottom connector on each controller is connected to Expander B.

The following figure shows the SAS cable connections for a two-controller-node configuration with two JBOD enclosures.

NOTE To save space, the figure does not show the disk drives that are in the JBOD enclosures.

Figure 6 Two-Controller-Node Configuration with Dual JBOD Enclosures

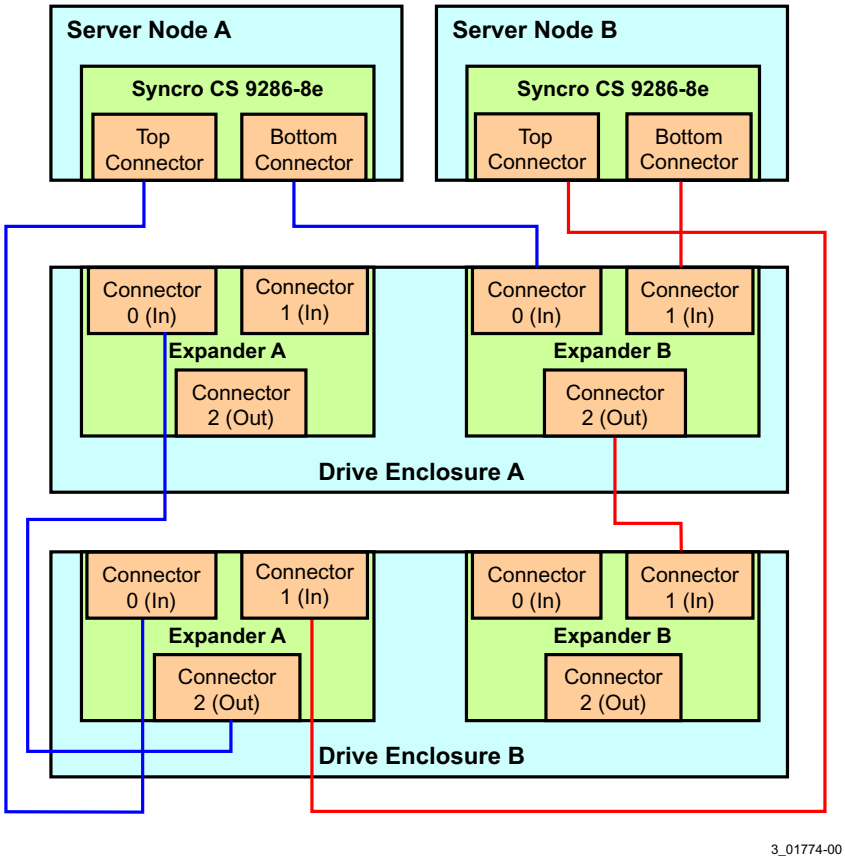

The recommended method shown in the preceding figure is preferable to simply daisy-chaining a second JBOD enclosure from the single-JBOD configuration shown in Figure 5, because a single power failure in the first JBOD enclosure could interrupt all data access. Instead, connect the second Syncro CS controller envisioning the JBOD enclosures in reverse order. The resulting *top-down/bottom-up* cabling approach shown in the preceding figure is preferred because it assures continued access to operating drives if either of the JBOD enclosures fails or is removed.

The following figure shows how to hot-add a second JBOD enclosure to an existing two-server cluster without interrupting service on the HA configuration.

NOTE To save space, the figure does not show the disk drives that are in the JBOD enclosures.

Figure 7 Adding a Second JBOD Enclosure - Redundant Configuration

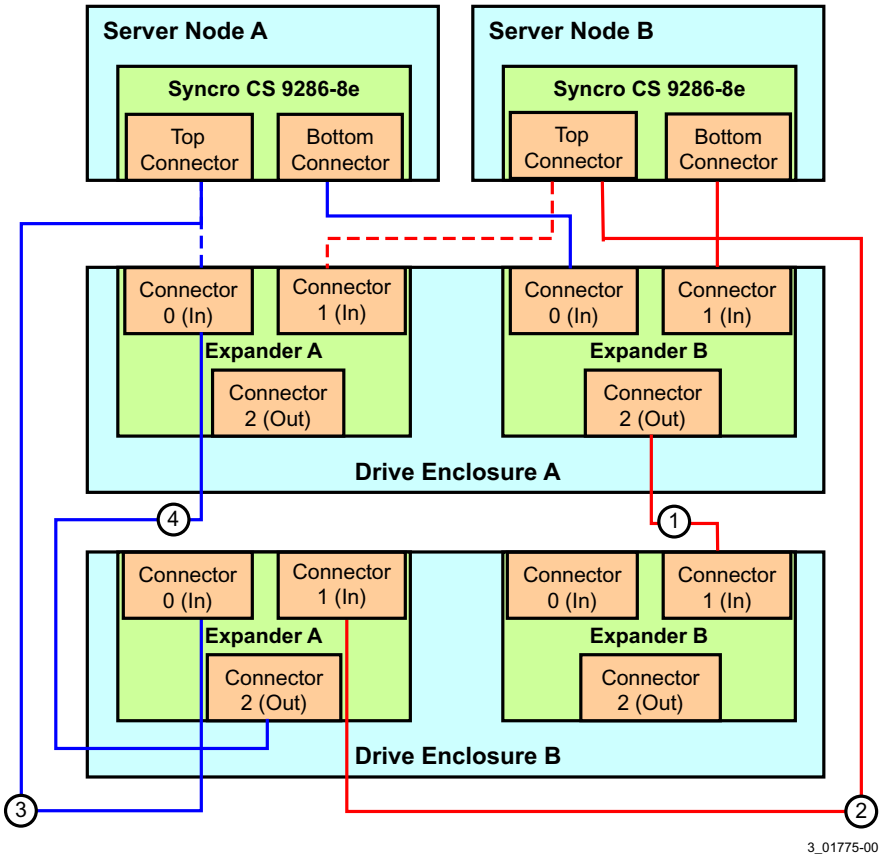

The steps for adding the second JBOD enclosure are as follows:

1. Connect a link from connector 2 on Expander B of JBOD enclosure A to connector 1 on expander B of JBOD enclosure B.
2. Disconnect the link from connector 1 on expander A of JBOD enclosure A and reconnect it to connector 1 on expander A of JBOD enclosure B.
3. Disconnect the link from connector 0 on expander A of JBOD enclosure A and reconnect it to connector 0 on expander A of JBOD enclosure B**.**
4. Connect the link from connector 2 on expander A of JBOD enclosure B to connector 0 on expander A of JBOD enclosure A.

Chapter 3: Creating the Syncro CS 9286-8e Configuration

This chapter explains how to set up HA-DAS clustering on a Syncro CS 9286-8e configuration after the hardware is fully configured and the operating system is installed.

3.1 Validating the Failover Configuration

Microsoft recommends that you validate the failover configuration before you set up failover clustering. To do this, run the Validate a Configuration wizard for Windows Server 2008 R2 or Windows Server 2012, following the instructions from Microsoft. The tests in the validation wizard include simulations of cluster actions. The tests fall into the following categories:

- **System Configuration tests**. These tests analyze whether the two server modules meet specific requirements, such as running the same version of the operating system version using the same software updates.
- **Network tests**. These tests analyze whether the planned cluster networks meet specific requirements, such as requirements for network redundancy.
- **Storage tests**. These tests analyze whether the storage meets specific requirements, such as whether the storage correctly supports the required SCSI commands and handles simulated cluster actions correctly.

Follow these steps to run the Validate a Configuration wizard.

NOTE You can also run the Validate a Configuration wizard after you create the cluster.

1. In the failover cluster snap-in, in the console tree, make sure Failover Cluster Management is selected and then, under Management, click **Validate a Configuration**.

   The Validate a Configuration wizard starts.
2. Follow the instructions for the wizard and run the tests.

   Microsoft recommends that you run all available tests in the wizard.

   NOTE Storage Spaces does not currently support Clustered RAID controllers. Therefore, do not include the *Validate Storage Spaces Persistent Reservation* storage test in the storage test suite. For additional information, visit the *Cluster in a Box Validation Kit for Windows Server* site on the Microsoft Windows Server TechCenter website.
3. When you arrive at the Summary page, click **View Reports** to view the results of the tests.
4. If any of the validation tests fails or results in a warning, correct the problems that were uncovered and run the test again.

3.2 Creating the Cluster

The Microsoft Server 2012 operating system installation does not enable the clustering feature by default. Follow these steps to view the system settings, and, if necessary, to enable clustering.

1. From the desktop, launch Server Manager.
2. Click **Manage** and select **Add Roles and Features**.
3. If the Introduction box is enabled (and appears), click **Next**.
4. In the Select Installation Type box, select **Role Based or Feature Based**.

5. In the Select Destination Server box, select the system and click **Next**.
6. In the Select Server Roles list, click **Next** to present the Features list.
7. Make sure that failover clustering is installed, including the tools. If necessary, run the Add Roles and Features wizard to install the features dynamically from this user interface.
8. If the cluster nodes need to support I/O as iSCSI targets, expand **File and Storage Services**, **File Services** and check for iSCSI Target Server and Server for NFS.

The Server Manager includes a configuration validator under **Server Manager> Tools> Failover Cluster Manager…Validate a Configuration**. Refer to the Microsoft documentation for more detailed information.

During creation of the cluster, Windows automatically defines and creates the *quorum*, a configuration database that contains metadata required for the operation of the cluster. To create a shared VD for the quorum, see the instructions in Section 3.3, Creating Virtual Drives on the Controller Nodes.

NOTE The recommended best practice is to create a small redundant VD for the quorum. A size of 500 MB is adequate for this purpose.

To determine if the cluster is active, run MSM and look at the **Dashboard** tab for the controller. The first of two nodes that boots shows the cluster status as *Inactive* until the second node is running and the MSM dashboard on the first node has been refreshed.

NOTE To refresh the MSM dashboard, press F5 or select **Manage > Refresh** on the menu.

The following figure shows the controller dashboard with *Active* peer controller status.

Figure 8 Controller Dashboard: Active Cluster Status

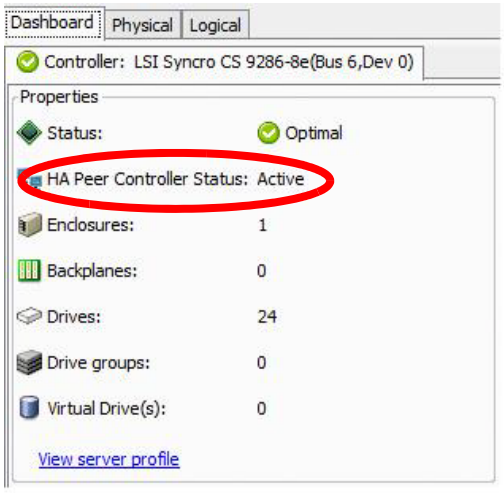

3.3 Creating Virtual Drives on the Controller Nodes

The next step is creating VDs on the disk drives.

The HA-DAS cluster configuration requires a minimum of one shared VD to be utilized as a quorum disk to enable Microsoft operating system support for clusters. Refer to the *MegaRAID SAS Software User Guide* for information about the available RAID levels and the advantages of each one.

As explained in the instructions in the following sections, VDs created for storage in an HA-DAS configuration must be shared. If you do not designate them as shared, the VDs are visible only from the controller node from which they were created.

You can use the WebBIOS pre-boot utility to create the VDs. You can also use the LSI MegaRAID Storage Manager (MSM) utility or the MegaCLI utility to create VDs after Windows has booted. Refer to the *MegaRAID SAS Software User Guide* for complete instructions on using these utilities.

3.3.1 Creating Shared or Exclusive VDs with the WebBIOS Utility

To coordinate the configuration of the two controller nodes, both nodes must be booted into the WebBIOS pre-boot utility. The two nodes in the cluster system boot simultaneously after power on, so you must rapidly access both consoles. One of the systems is used to create the VDs; the other system simply remains in the pre-boot utility. This approach keeps the second system in a state that does not fail over while the VDs are being created on the first system.

NOTE The WebBIOS utility cannot see boot sectors on the disks. Therefore, be careful not to select the boot disk for a VD. Preferably, unshare the boot disk before doing any configuration with the pre-boot utility. To do this, select **Logical Drive Properties** and deselect the **Shared Virtual Disk** property.

Follow these steps to create VDs with the WebBIOS utility.

1. When prompted during the POST on the two systems, use the keyboard to access the WebBIOS pre-boot BIOS utility (on both systems) by pressing **Ctrl-H**.

   Respond quickly, because the system boot times are very similar and the time-out period is short. When both controller nodes are running the WebBIOS utility, follow these steps to create RAID 5 arrays.

   NOTE To create a RAID 0, RAID 1, or RAID 6 array, modify the instructions to select the appropriate number of disks.

2. Click **Start**.
3. On the WebBIOS main page, click **Configuration Wizard**, as shown in the following figure.

Figure 9 WebBIOS Main Page

The first Configuration Wizard window appears.

4. Select **Add Configuration** and click **Next**.
5. On the next wizard screen, select **Manual Configuration** and click **Next**.

   The Drive Group Definition window appears.
6. In the Drives panel on the left, select the first drive, then hold down the Ctrl key and select more drives for the array, as shown in the following figure.

Figure 10 Selecting Drives

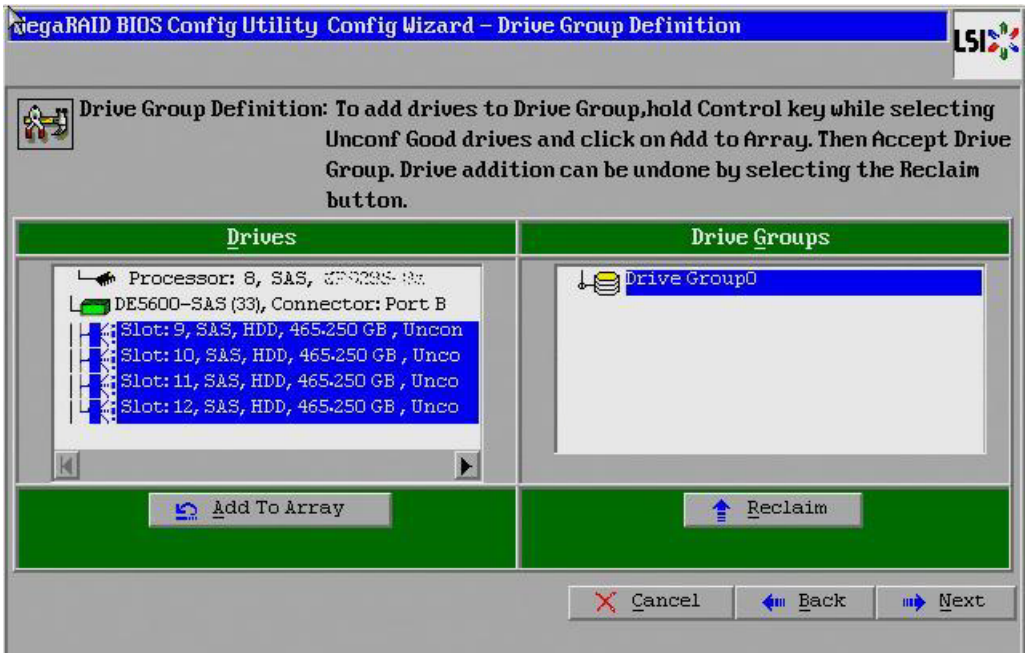

7. Click **Add To Array**, click **ACCEPT**, and click **Next**.
8. On the next screen, click **Add to SPAN**, then click **Next**.
9. On the next screen, click **Update Size**.

10. Select **Provide Shared Access** on the bottom left of the window, as shown in the following figure.

    Alternatively, deselect this option to create an exclusive VD as a boot volume for this cluster node.

Figure 11 Virtual Drive Definition

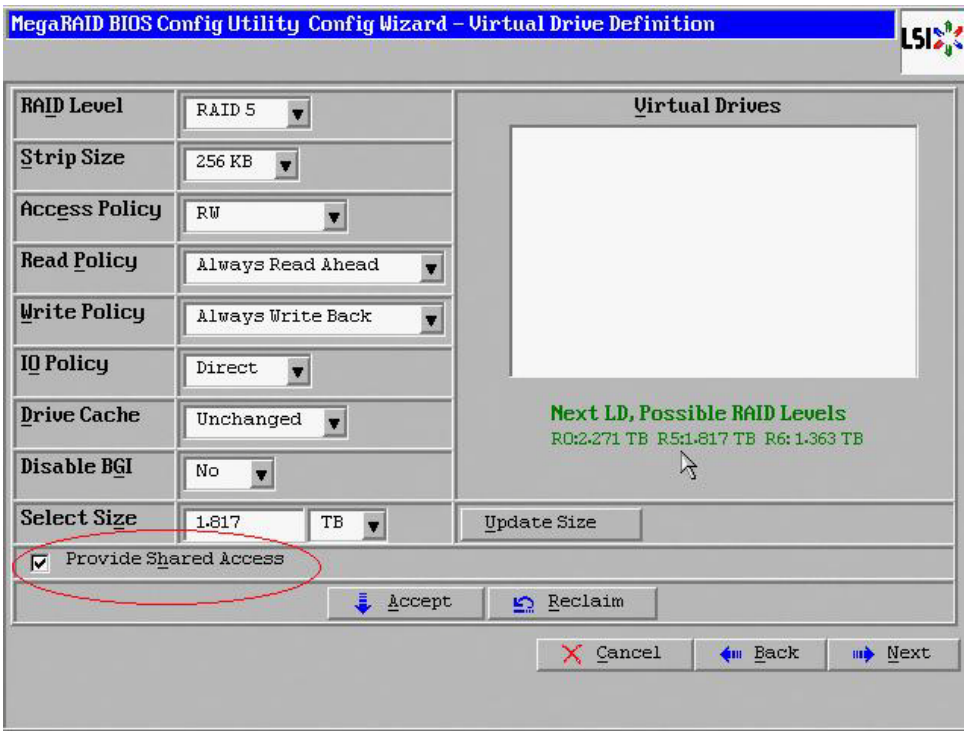

The **Provide Shared Access** option enables a shared VD that both controller nodes can access. If you uncheck this box, the VD has a status of *Exclusive*, and only the controller node that created this VD can access it.

11. On this same page, click **Accept**, then click **Next**.
12. On the next page, click **Next**.
13. Click **Yes** to accept the configuration.
14. Repeat the previous steps to create the other VDs.

    As the VDs are configured on the first controller node, the other controller node's drive listing is updated to reflect the use of the drives.
15. When prompted, click **Yes** to save the configuration, and click **Yes** to confirm that you want to initialize it.
16. Define hot spare disks for the VDs to maximize the level of data protection.

    NOTE The Syncro CS 9286-8e solution supports global hot spares and dedicated hot spares. Global hot spares are global for the cluster, not for a controller.
17. When all VDs are configured, reboot both systems as a cluster.

3.3.2 Creating Shared or Exclusive VDs with MegaCLI64.exe on Windows Server 2012

MegaCLI is a command-line-driven utility used to create and manage VDs. MegaCLI can run in any directory on the server. The following procedure assumes that a current copy of the 64-bit Windows version of MegaCLI is located on the server in `c:\lsi\cli`.

The steps for creating a VD are slightly different depending on whether you run MegaCLI in Windows PowerShell® or from a Windows command prompt. Therefore, two sets of instructions are included.

3.3.2.1 Creating Shared or Exclusive VDs: Running MegaCLI in Windows PowerShell

Follow these steps to create a shared VD with MegaCLI on Microsoft Windows Server 2012 running in Windows PowerShell.

NOTE Enter the command line entries exactly as shown in the following instructions, because the syntax used for PowerShell is slightly different than the syntax used for the Windows command prompt. The MegaCLI help and documentation do *not* include syntactical references to PowerShell.

1. On the Microsoft Server 2012 desktop, right-click the PowerShell icon and select **Run as Administrator** from the pop-up menu, as shown in the following figure.

Figure 12 Starting PowerShell

2. At the PowerShell prompt, enter the command `cd \lsi\cli` to change to the MegaCLI directory.
3. At the PowerShell prompt, run the following command:

   ```
   .\megacli64 "-cfgdsply -a0"
   ```

   The `-a0` parameter presumes that there is only one Syncro CS 9286-8e controller in the system or that these steps reference the first Syncro CS 9286-8e controller in a system with multiple controllers.

The following figure shows some sample configuration information that appears in response to the command.

Figure 13 Sample Configuration Information

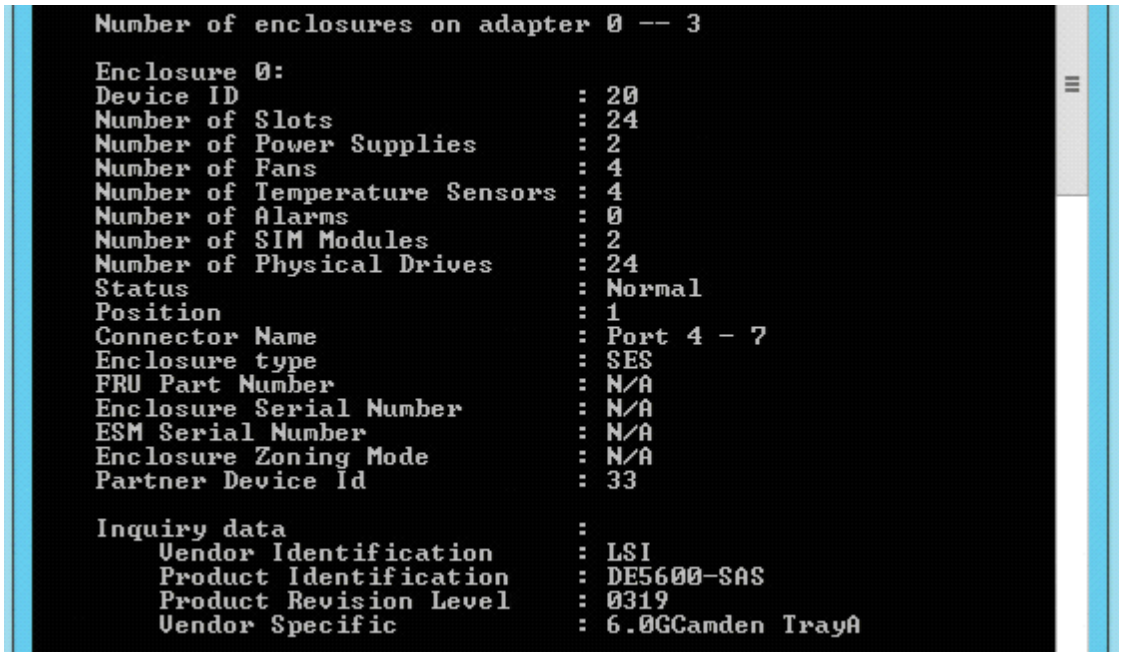

The command generates many lines of information that scroll down in the PowerShell window. You need to use some of this information in the command line to create the shared VD.

4. Find the JBOD enclosure number Device ID for the system and also the Device IDs of available physical drives for the VD you will create.

   In the preceding figure, the enclosure device ID of **20** appears close to the top of the window. Use the scroll bar to find the device IDs of the physical drives for the VD. The format of the drive IDs appears as follows:

   ```
   Physical Disk: 1
   Enclosure Device ID: 20
   Slot Number: 2
   Drive's position: DiskGroup: 0, Span; 0, Arm: 1
   Enclosure position: 0
   Device Id: 1
   WWN: 5000C5001AA6F4DC
   ```

5. Create the shared VD using the enclosure and disk device IDs with the following command line syntax:

   ```
   .\megacli64 "-cfgldadd -r5[20:1,20:2,20:3,20:4,20:5] WB RA direct -strpsz64 -a0"
   ```

   The following notes explain the command line parameters.

   — The `-cfgldadd` parameter configures and adds a VD (logical disk).
   — The `-r5` parameter selects RAID 5 as the RAID level.
   — The opening and closing square brackets define the list of drives for the VD. Each drive is listed in the form *enclosure device ID*: *drive device ID*.
   — The `WB` parameter sets the controller to use the write cache. (Alternatively, the `WT` parameter sets the controller cache for write through.)
   — The RA parameter sets the controller cache for read ahead.
   — The `direct` parameter sets direct I/O.
   — The `-strpsz64` parameter sets the stripe size to 64 KB.
   — The `-a0` parameter selects the first Syncro CS 9286-8e controller in the system.

— The HA-DAS version of MegaCLI creates, by default, a *shared* VD that is visible to all cluster nodes.

NOTE To create a VD that is visible only to the node that created it (such as creating a boot volume for this cluster node), add the `-exclusive` parameter to the command line.

— The `-a0` parameter selects the first Syncro CS 9286-8e controller in the system.

For more information about MegaCLI command line parameters, refer to the *MegaRAID SAS Software User Guide*.

3.3.2.2 Creating Shared or Exclusive VDs: Running MegaCLI from a Windows Command Prompt

Follow these steps to create shared VDs with MegaCLI on Microsoft Windows Server 2012, running in a Windows command prompt:

1. Move the mouse to the lower right hand corner of the screen and select Search when the icons appear, as shown in the following figure.

Figure 14 Selecting Search

2. Search for **cmd**, as shown in the following figure.

Figure 15 Searching for 'cmd'

The system finds the command prompt.

3. Right-click on the Command Prompt. A properties check indicator appears on the button, as shown in the following figure.

Figure 16 Command Prompt

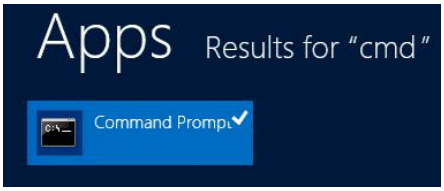

The following button appears at the bottom of the desktop to open the command prompt as Administrator.

Figure 17 Run as Administrator Button

4. Click the button to open a command prompt, as shown in the following figure.

Figure 18 Windows Command Prompt

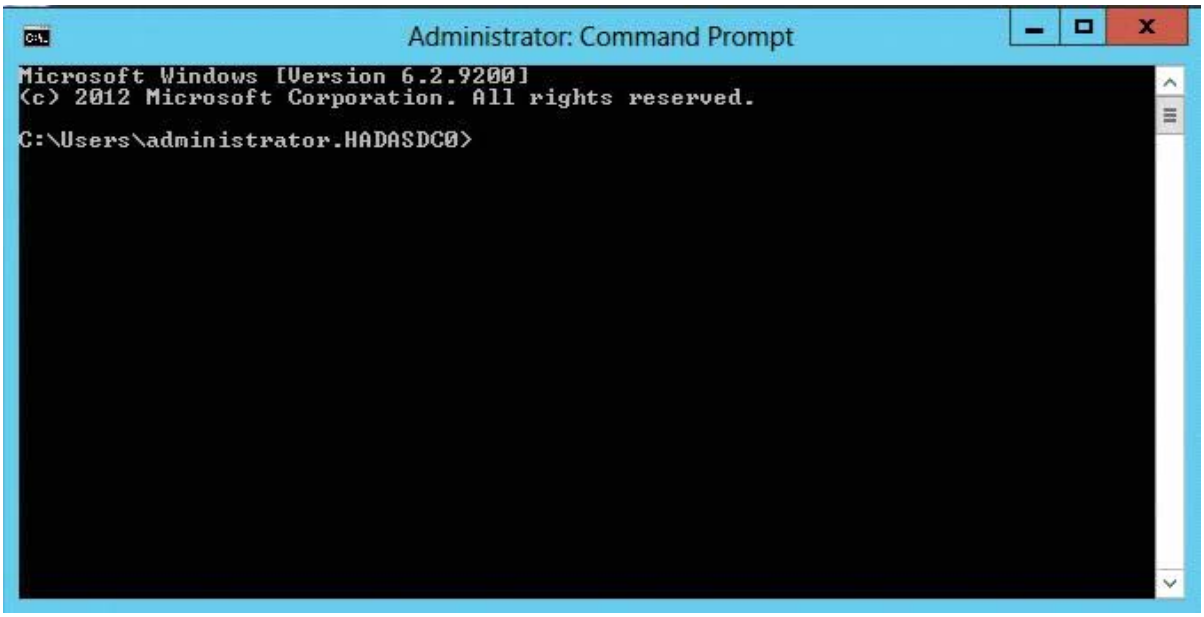

5. Enter the following command at the prompt:

```
megacli64 –cfgdsply –a0
```

The `-a0` parameter presumes that there is only one Syncro CS 9286-8e controller in the system or that these steps reference the first Syncro CS 9286-8e controller in a system with multiple controllers.

The following figure shows some sample configuration information that appears in response to the command.

Figure 19 Sample Configuration Information

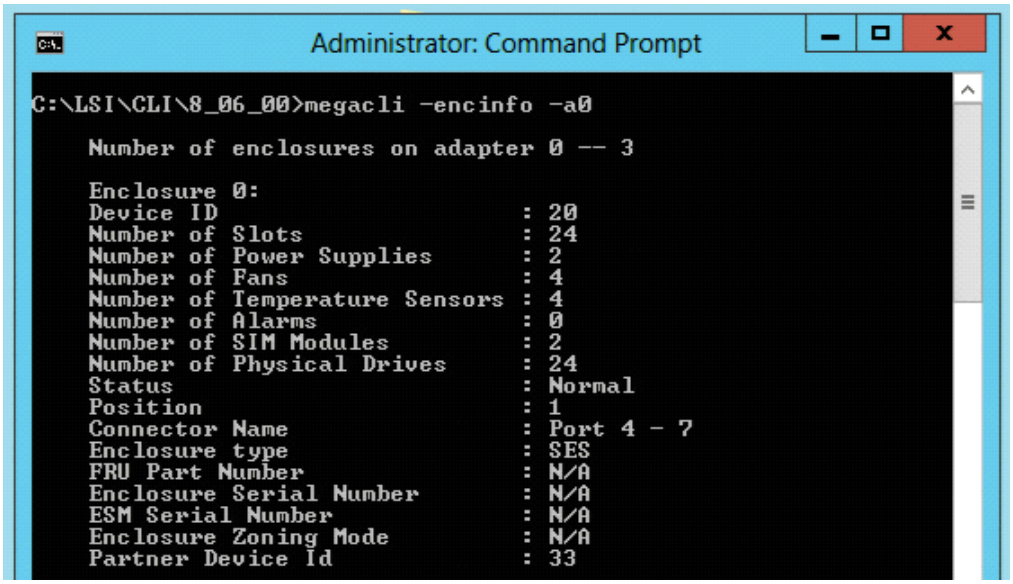

The command generates many lines of information that scroll down in the command prompt window. You need to use some of this information in the command line to create the shared VD.

6. Find the JBOD enclosure number Device ID for the system and also the Device IDs of available physical drives for the VD you will create.

   In the preceding figure, the enclosure device ID of **20** appears close to the top of the window. Use the scroll bar to find the device IDs of the physical drives for the VD. The format of the drive IDs appears as follows:

   ```
   Physical Disk: 1
   Enclosure Device ID: 20
   Slot Number: 2
   Drive’ s position: DiskGroup: 0, Span; 0, Arm: 1
   Enclosure position: 0
   Device Id: 1
   WWN: 5000C5001AA6F4DC
   ```

7. Create the shared VD using the enclosure and disk device IDs with the following command line syntax:

   ```
   megacli64 –cfgldadd –r5[20:1,20:2,20:3,20:4,20:5] WB RA direct –strpsz64 –a0
   ```

   The following notes explain the command line parameters.

   — The `–cfgldadd` parameter configures and adds a VD (logical disk).
   — The `-r5` parameter selects RAID 5 as the RAID level.
   — The opening and closing square brackets define the list of drives for the VD. Each drive is listed in the form *enclosure device ID*: *drive device ID*.
   — The `WB` parameter sets the controller to use the write cache. (Alternatively, the `WT` parameter sets the controller cache for write through.)
   — The RA parameter sets the controller cache for read ahead.
   — The `direct` parameter sets direct I/O.
   — The `–strpsz64` parameter sets the stripe size to 64 KB.
   — The `-a0` parameter selects the first Syncro CS 9286-8e controller in the system.
   — The HA-DAS version of MegaCLI creates, by default, a *shared* VD that is visible to all cluster nodes.

NOTE To create a VD that is visible only to the node that created it (such as creating a boot volume for this cluster node), add the `-exclusive` parameter to the command line.

— The `-a0` parameter selects the first Syncro CS 9286-8e controller in the system.

For more information about the MegaCLI command line parameters, refer to the *MegaRAID SAS Software User Guide*.

3.3.3 Creating Shared or Exclusive VDs with MSM

Follow these steps to create VDs for data storage with MSM. When you create the VDs, you assign the Share Virtual Drive property to them to make them visible from both controller nodes. This example assumes you are creating a RAID 5 redundant VD. Modify the instructions as needed for other RAID levels.

NOTE Not all versions of MSM support HA-DAS. Check the release notes to determine if your version of MSM supports HA-DAS. Also, see Section 5.1, Verifying HA-DAS Support in Tools and the OS Driver.

1. In the left panel of the MSM Logical pane, right-click the Syncro CS 9286-8e controller and select **Create Virtual Drive** from the pop-up menu.

   The Create Virtual Drive wizard appears.
2. Select the **Advanced** option and click **Next**.
3. In the next wizard screen, select **RAID 5** as the RAID level, and select unconfigured drives for the VD, as shown in the following figure.

Figure 20 Drive Group Settings

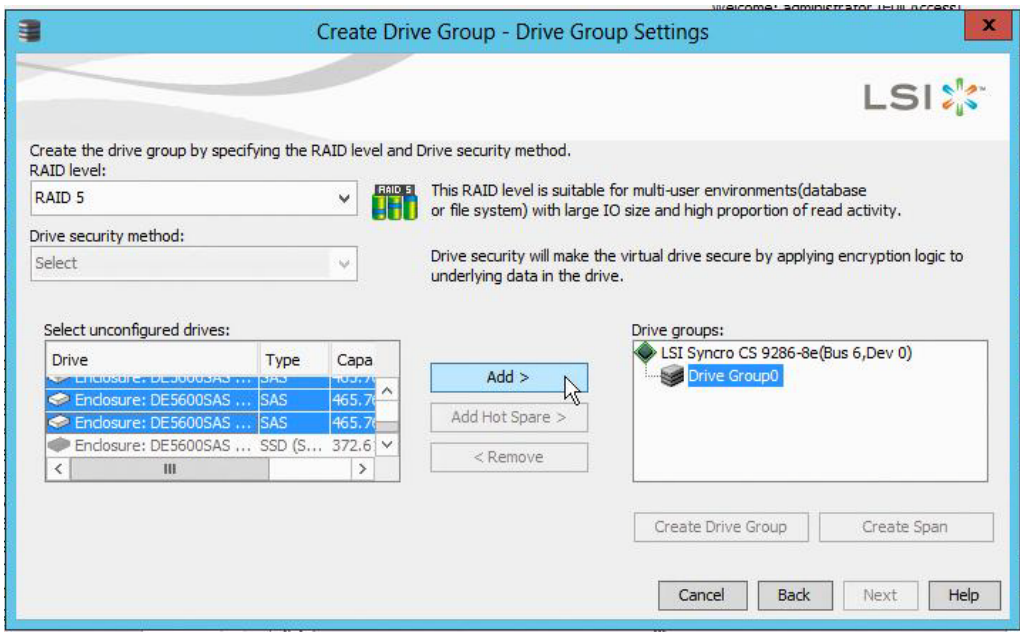

4. Click **Add** to add the VD to the drive group.

   The selected drives appear in the Drive groups window on the right.
5. Click **Create Drive Group**. Then click **Next** to continue to the next window.

   The Virtual Drive Settings window appears.
6. Enter a name for the VD.

7. Select **Always Write Back** as the Write policy option, and select other VD settings as required.
8. Select the **Provide Shared Access** option, as shown in the following figure.

NOTE If you do not select **Provide Shared Access**, the VD is visible only from the server node on which it is created. Leave this option unselected if you are creating a boot volume for this cluster node.

Figure 21 Provide Shared Access Option

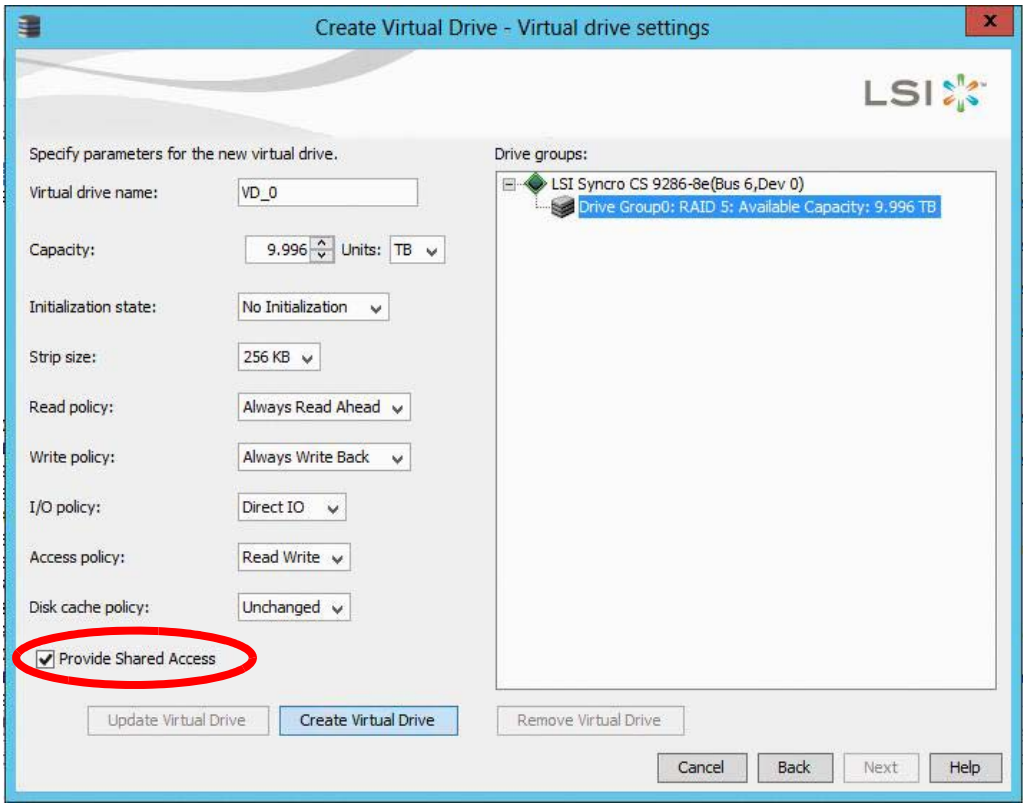

9. Click **Create Virtual Drive** to create the virtual drive with the settings you specified. The new VD appears in the Drive groups window on the right of the window.
10. Click **Next** to continue.

The Create Virtual Drive Summary window appears, as shown in the following figure.

Figure 22 Create Virtual Drive Summary

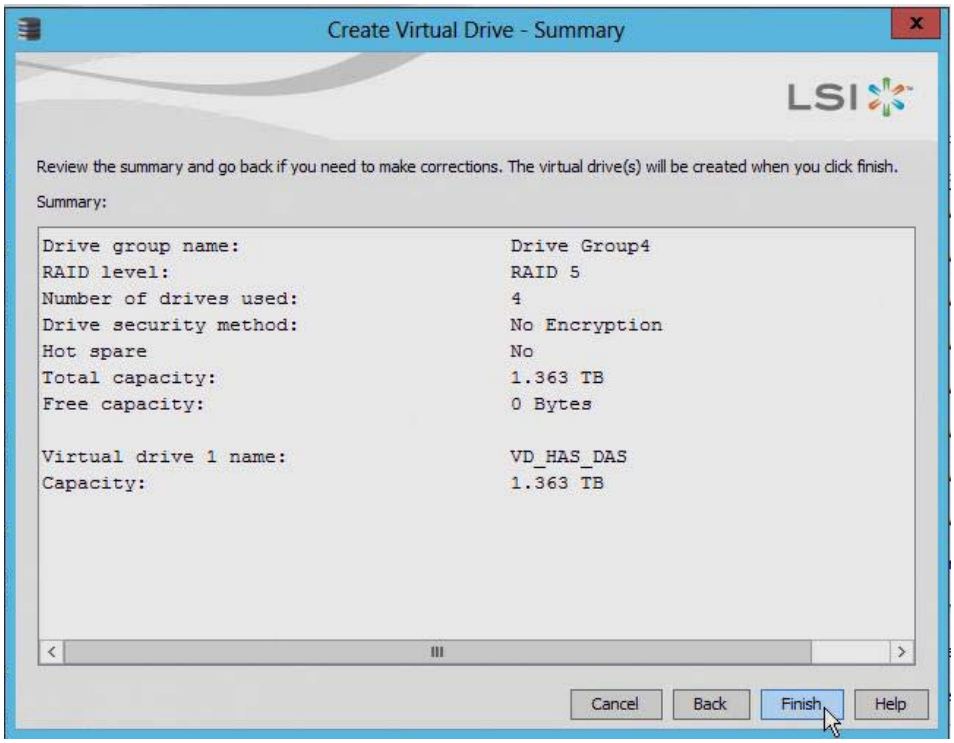

11. Click **Finish** to complete the VD creation process.
12. Click **OK** when the Create Virtual Drive - complete message appears.

3.3.3.1 Unsupported Drives

Drives that are used in the Syncro CS 9286-8e solution must selected from the list of approved drives listed on the LSI web site (see the URL in 1.4, Hardware Compatibility). If the MegaRAID Storage Manager (MSM) utility finds a drive that does not meet this requirement, it marks the drive as *Unsupported*, as shown in the following figure.

Figure 23 Unsupported Drive in MSM

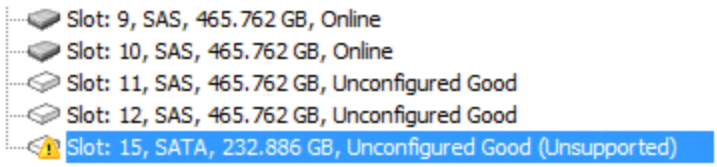

3.4 HA-DAS CacheCade Support

The Syncro CS 9286-8e controller includes support for CacheCade 1.0, a feature that uses SAS SSD devices for read caching of frequently accessed read data. When a VD is enabled for the CacheCade feature, frequently read data regions of the VD are copied into the SSD when the CacheCade algorithm determines the region is a good candidate. When the data region is in the CacheCade SSD volume, the firmware can service related reads from the faster-access SSD volume instead of the higher-latency VD. The CacheCade feature uses a single SSD to service multiple VDs.

The Syncro CS 9286-8e solution requires the use of SAS SSDs that support SCSI-3 persistent reservations (PR) for CacheCade VDs. LSI maintains a list of SAS SSD drives that meet the HA-DAS requirements

NOTE A CacheCade VD is not presented to the host operating system, and it does not move to the peer controller node when a failover occurs. A CacheCade VD possesses properties that are similar to a VD with exclusive host access. Therefore, the CacheCade volume does not cache read I/Os for VDs that are managed by the peer controller node.

Follow these steps to create a CacheCade 1.0 VD as part of a Syncro CS 9286-8e configuration. The procedure automatically associates the CacheCade volume with all existing shared VDs in the configuration. Be sure that one or more SAS SSD drives are already installed in the system. Also, be sure you are using a version of MSM that supports HA-DAS.

1. In MSM, open the physical view, right-click on the controller name and select **Create CacheCade SSD Caching**.
2. In the Drive Group window, set the CacheCade RAID level and select one or more unconfigured SSD drives. Use the **Add** button to place the selected drives into the drive group.

   RAID 0 is the recommended RAID level for the CacheCade volume.

   The following figure shows the CacheCade drive group.

Figure 24 Creating a CacheCade Drive Group: 1

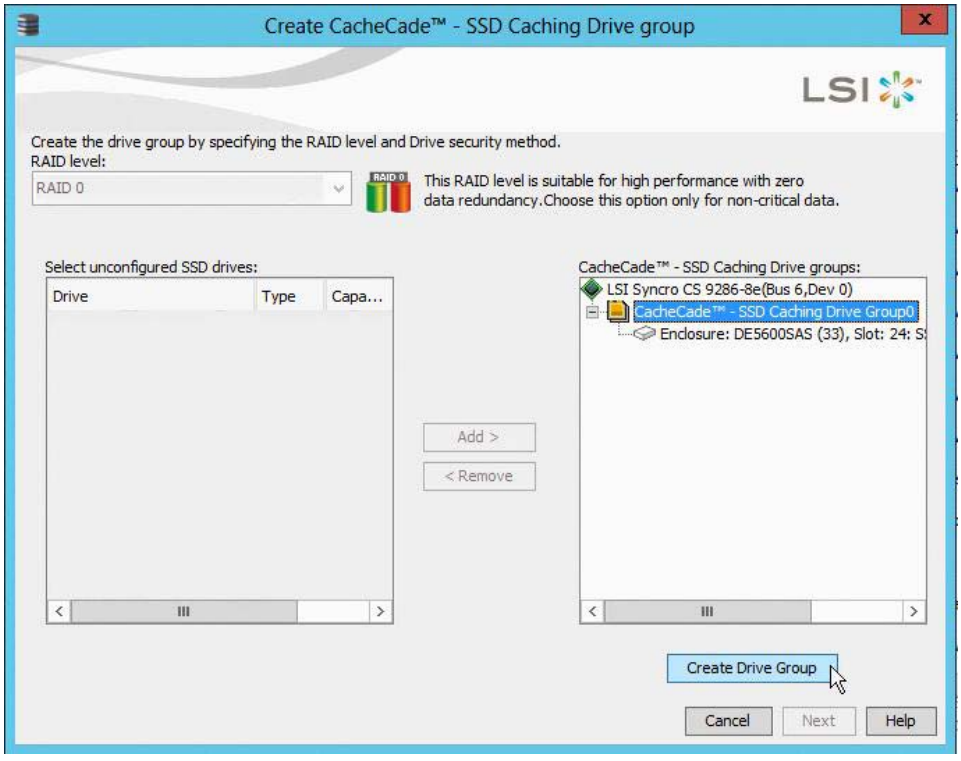

3. Click **Create Drive Group** and then click **Next**.
4. In the Create CacheCade SSD Caching Virtual Drive window, update the SSD Caching VD name and set the size as necessary.

   The maximum allowable size for the CacheCade volume is 512 GB. To achieve optimal read cache performance, the recommended best practice is to make the size as large as possible with the available SSDs, up to this limit.

Figure 25 Creating a CacheCade Drive Group: 2

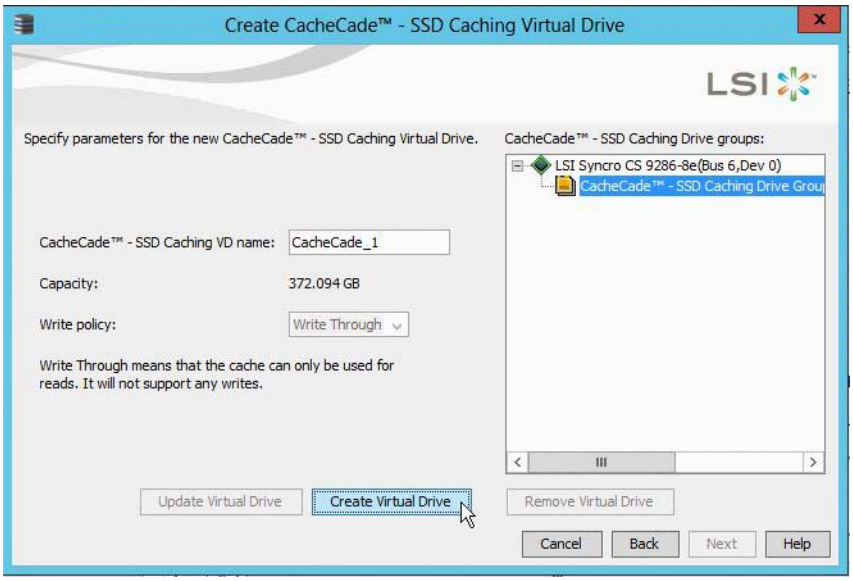

5. Click **Create Virtual Drive** and then click **Next**.
6. In the Create CacheCade SSD Caching Summary window, review the configuration and then click **Finish**.

Figure 26 Reviewing the Configuration

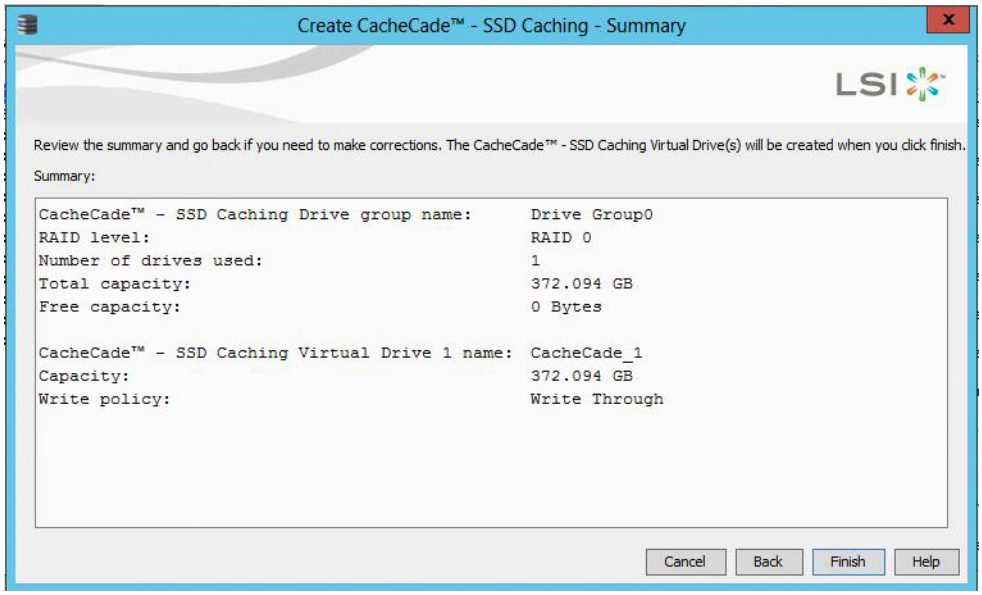

7. In the Create CacheCade SSD Caching Complete box, click **OK**.

The CacheCade VD now appears on the Logical tab of MSM, as shown in the following figure. The CacheCade volume association with the drive groups appears in this view.

Figure 27 New CacheCade Drive Group

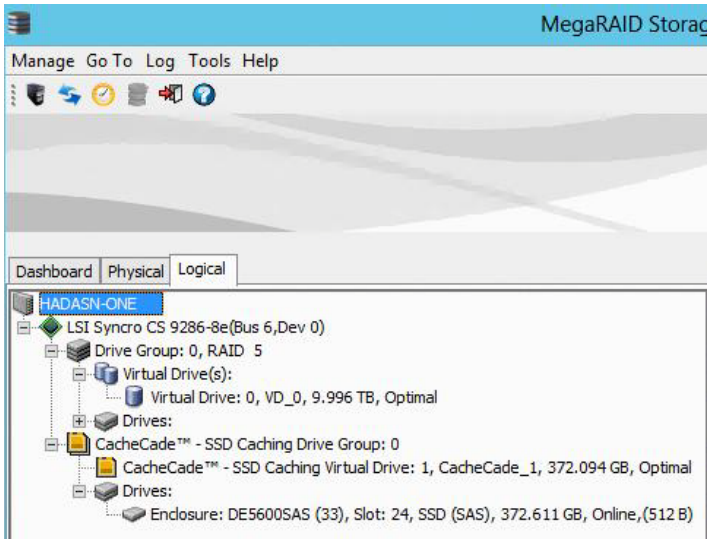

Chapter 4: System Administration

This chapter explains how to perform system administration tasks, such as planned failovers and updates of the Syncro CS 9286-8e controller firmware.

4.1 High Availability Properties

The following figure shows the high availability properties that MSM displays on the Controller Properties tab for a Syncro CS 9286-8e controller.

Figure 28 Controller Properties: High Availability Properties

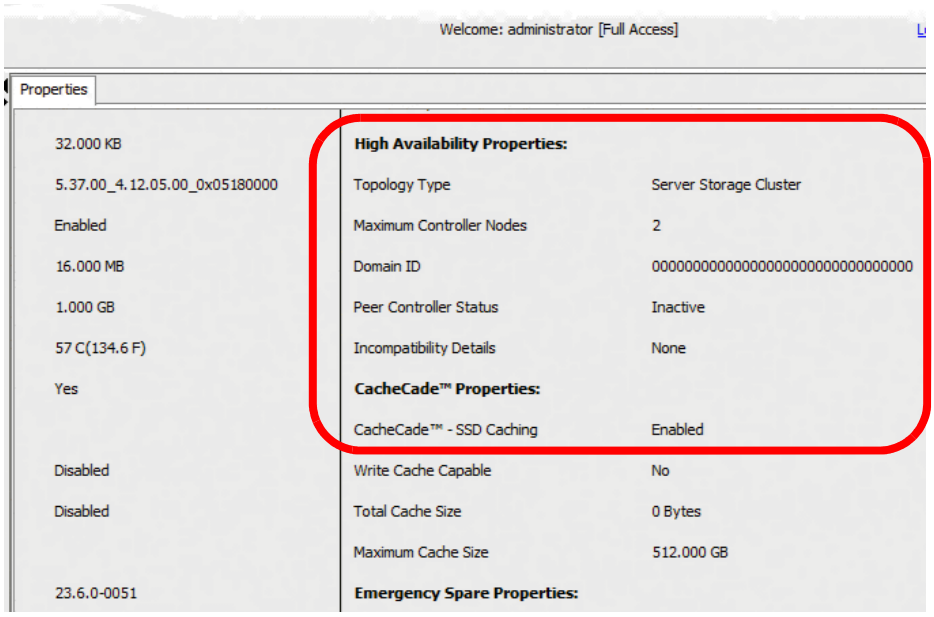

Following is a description of each high availability property:

- **Topology Type** – A descriptor of the HA topology for which the Syncro CS 9286-8e controller is currently configured (the default is *Server Storage Cluster*).
- **Maximum Controller Nodes** – The maximum number of concurrent Syncro CS 9286-8e controllers within the HA domain that the controller supports.
- **Domain ID** – A unique number that identifies the HA domain in which the controller is currently included. This field has a number if the cluster or peer controller is in active state.
- **Peer Controller Status** – The current state of the peer controller.
  Active: The peer controller is present and is participating in the HA domain.
  Inactive: The peer controller is missing or has failed.
  Incompatible: The peer controller is detected, but it has an incompatibility with the controller.
- **Incompatibility Details** – If the peer controller is incompatible, this field lists the cause of the incompatibility.

4.2 Understanding Failover Operations

A *failover operation* in HA-DAS is the process by which VD management transitions from one server node to the peer server node. A failover operation might result from a user-initiated, planned action to move an application to a different controller node so that maintenance activities can be performed, or the failover might be unintended and unplanned, resulting from hardware component failure that blocks access to the storage devices. Figure 29 and Figure 30 show an example of a failover operation of various drive groups and VDs from Server A to Server B. The following figure shows the condition of the two server nodes before the failover.

Figure 29 Before Failover from Server A to Server B

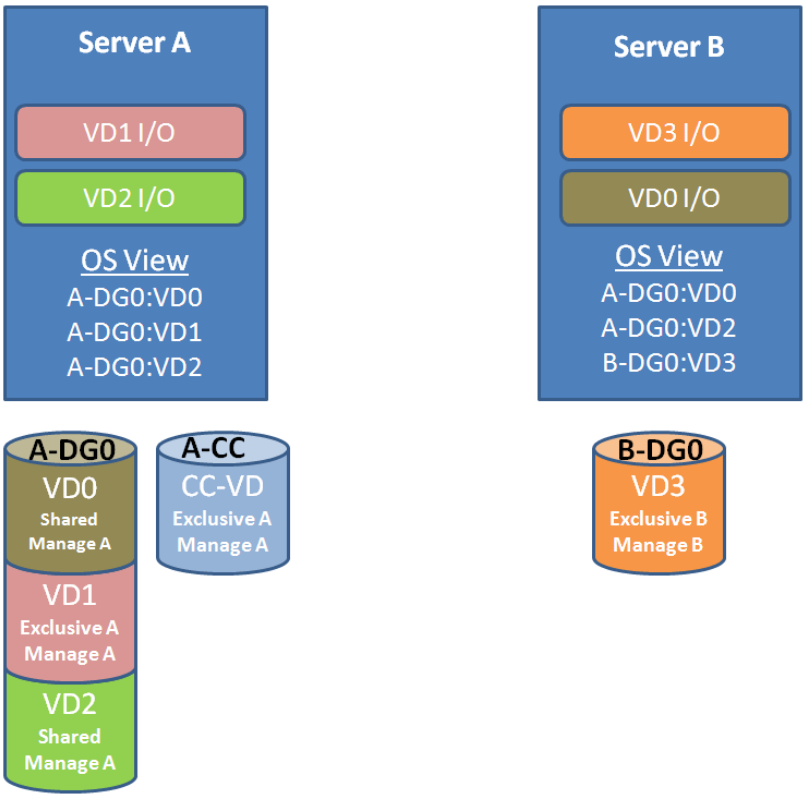

Before failover, the cluster status is as follows in terms of managing the drive group and VDs:

- All VDs in A-DG0 (Server A - Drive Group 0) are managed by Server A.
- VD3 in B-DG0 (Server B – Drive Group 0) is managed by Server B.
- The CacheCade VD (CC-VD) in A-CC is managed by Server A and services VDs in drive group A-DG0.

Before failover, the operating system perspective is as follows:

- The operating system on Server A only sees VDs with shared host access and exclusive host access to Server A.
- The operating system on Server B only sees VDs with shared host access and exclusive host access to Server B.

Before failover, the operating system perspective of I/O transactions is as follows:

- Server A is handling I/O transactions that rely on A-DG0:VD1 and A-DG0:VD2.
- Server B is handling I/O transactions that rely on A-DG0:VD0 and B-DG0:VD3.

The following figure shows the condition of the two server nodes after the failover.

Figure 30 After Failover from Server A to Server B

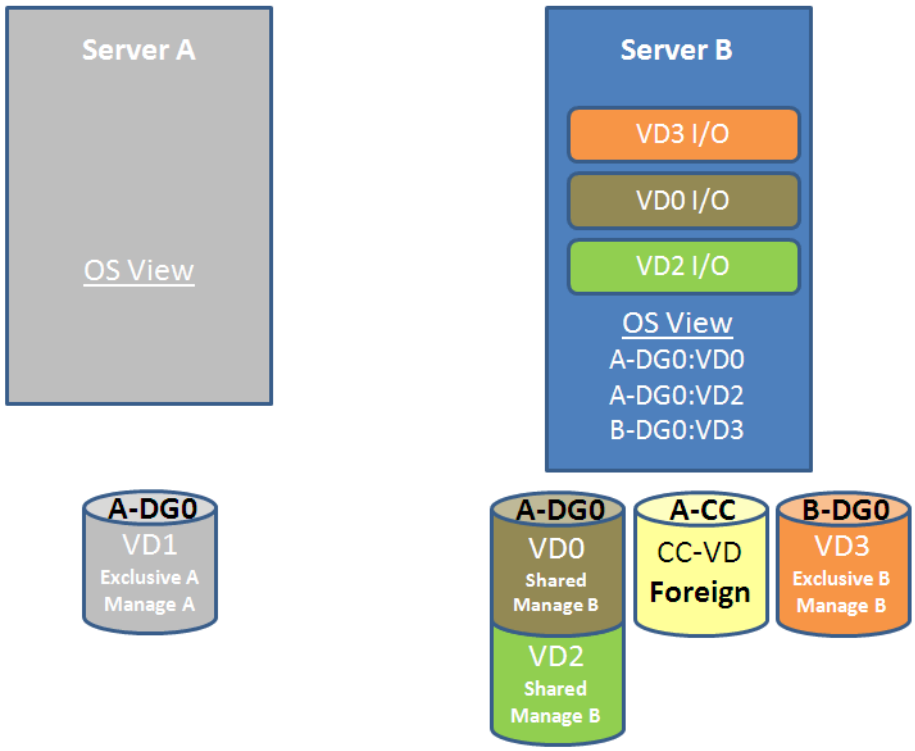

After failover, the cluster status is as follows, in terms of managing the drive group and the VDs:

- All shared VDs in A-DG0 have failed over and are now managed by Server B.
- VD3 in B-DG0 is still managed by Server B.
- The CacheCade VD (CC-VD) in A-CC now appears as a foreign VD on Server B but does not service any VDs in A-DG0 or B-DG0.

After failover, the operating system perspective is as follows:

- The operating system on Server B manages all shared VDs and any exclusive Server B VDs.

After failover, the operating system perspective of I/O transactions is as follows:

- Failover Cluster Manager has moved the I/O transactions for VD2 on A-DG0 to Server B.
- Server B continues to run I/O transactions on B-DG0:VD3.
- I/O transactions that rely on the exclusive A-DG0:VD1 on Server A fail because exclusive volumes do not move with a failover

NOTE When Server A returns, the management and I/O paths of the pre-failover configurations are automatically restored.

The following sections provide more detailed information about planned failover and unplanned failover.

4.2.1 Understanding and Using Planned Failover

A planned failover occurs when you deliberately transfer control of the drive groups from one controller node to the other. The usual reason for initiating a planned failover is to perform some kind of maintenance or upgrade on one of the controller nodes—for example, upgrading the controller firmware, as described in the following section. A planned failover can occur when there is active data access to the shared drive groups.

Before you start a planned failover on a Syncro CS system, be sure that no processes are scheduled to run during that time. Be aware that system performance might be impacted during the planned failover.

NOTE Failed-over VDs with exclusive host access cannot be accessed unless the VD host access is set to SHARED. Do not transition operating system boot volumes from EXCLUSIVE to SHARED access.

4.2.1.1 Planned Failover in Windows Server 2012

Follow these steps to perform a planned failover on a Syncro CS 9286-8e system running Windows Server 2012.

1. Create a backup of the data on the Syncro CS 9286-8e system.
2. In the Failover Cluster Manager snap-in, if the cluster that you want to manage is not displayed in the console tree, right-click **Failover Cluster Manager**, click **Manage a Cluster**, and then select or specify the cluster that you want.
3. If the console tree is collapsed, expand the tree under the cluster that you want to configure.
4. Expand **Services and Applications**, and click the name of the virtual machine.
5. On the right-hand side of the screen, under **Actions**, click **Move this service or application to another node**, and click the name of the other node.

As the virtual machine is moved, the status is displayed in the results panel (center panel). Verify that the move succeeded by inspecting the details of each node in the RAID management utility.

4.2.1.2 Planned Failover in Windows Server 2008 R2

Follow these steps to perform a planned failover on a Syncro CS 9286-8e system running Windows Server 2008 R2.

1. Create a backup of the data on the Syncro CS 9286-8e system.
2. Open the Failover Cluster Manager, as shown in the following figure.

Figure 31 Failover Cluster Manager

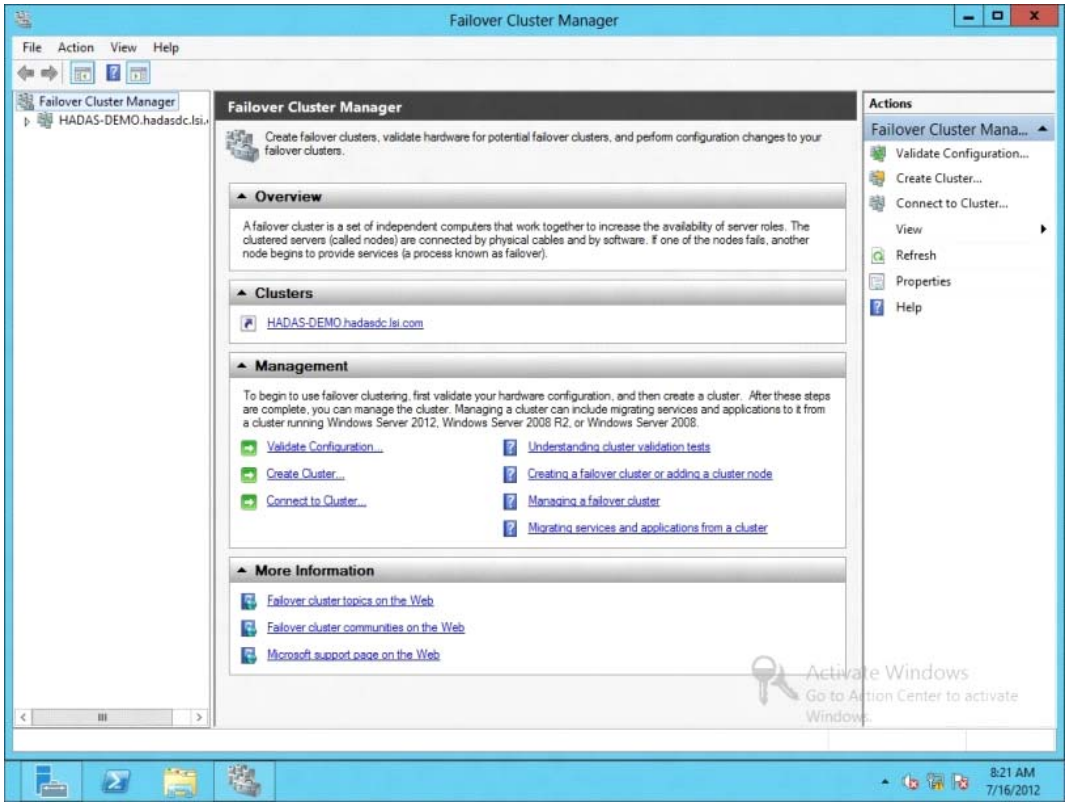

3. In the left panel, expand the tree to display the disks, as shown in the following figure.

Figure 32 Expand Tree

4. Right-click on the entry in the Assigned To column in the center panel of the window.
5. On the pop-up menu, select **Move** > **Select Node**, as shown in the following figure.

Figure 33 Expand Tree

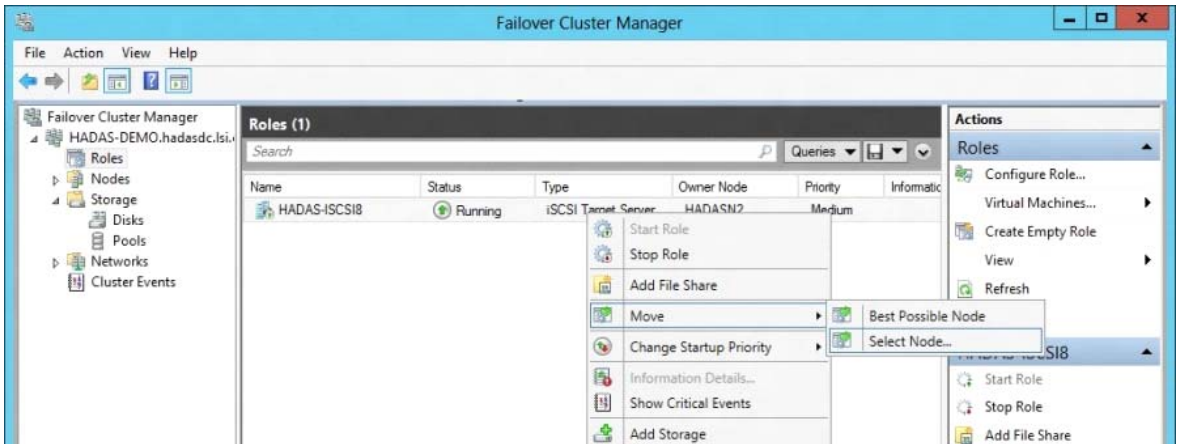

6. Select the node for the planned failover.

4.2.2 Understanding Unplanned Failover

An unplanned failover might occur if the controller in one of the server nodes fails, or if the cable from one controller node to the JBOD enclosure is accidentally disconnected. The Syncro CS 9286-8e solution is designed to automatically switch to the other controller node when such an event occurs, without any disruption of access to the data on the drive groups.

NOTE When the failed controller node returns, the management and I/O paths of the pre-failover configurations are automatically restored.

4.3 Updating the Syncro CS 9286-8e Controller Firmware

Follow these steps to update the firmware on the Syncro CS 9286-8e controller board. You must perform the update only on the controller node that is *not* currently accessing the drive groups.

NOTE Be sure that the version of firmware selected for the update is specified for Syncro CS controllers. If you updating to a version of controller firmware that does not support Syncro CS controllers, you will experience a loss of HA-DAS functionality.

1. If necessary, perform a planned failover as described in the previous section to transfer control of the drive groups to the other controller node.
2. Start the MSM utility on the controller node that does *not* currently own the cluster.

NOTE To determine which node currently owns the cluster in Windows Server 2012, follow the steps in Section 4.2.1.2, Planned Failover in Windows Server 2008 R2, up to step 3, where information about the cluster disks is displayed in the center panel. The current owner of the cluster is listed in the Owner Node column.

3. In the left panel of the MSM window, click the icon of the controller that requires an upgrade.
4. In the MSM window, select **Go To > Controller > Update Controller Firmware**.
5. Click **Browse** to locate the `.rom` update file.

6. After you locate the file, click **Ok**.
   The MSM software displays the version of the existing firmware and the version of the new firmware file.
7. When you are prompted to indicate whether you want to upgrade the firmware, click **Yes**.
   The controller is updated with the new firmware code contained in the `.rom` file.
8. Reboot the controller node after the new firmware is flashed.
   The new firmware does not take effect until reboot.
9. If desired, use planned failover to transfer control of the drive groups back to the controller node you just upgraded.
10. Repeat this process for the other controller.
11. Restore the cluster to its non-failed-over mode.

4.4 Updating the MegaRAID Driver

To update the MegaRAID driver used in the clustering configuration, download the latest version of the driver from the LSI website. Then follow these instructions for Windows Server 2008 R2 or Windows Server 2012.

4.4.1 Updating the MegaRAID Driver in Windows Server 2008 R2

As a recommended best practice, always back up system data before updating the driver, and then perform a planned failover. These steps are recommended because a driver update requires a system reboot.

1. Right-click on **Computer** and select **Properties**.
2. Click **Change Settings**, as shown in the following figure.

Figure 34 Windows Server 2008 R2 System Properties

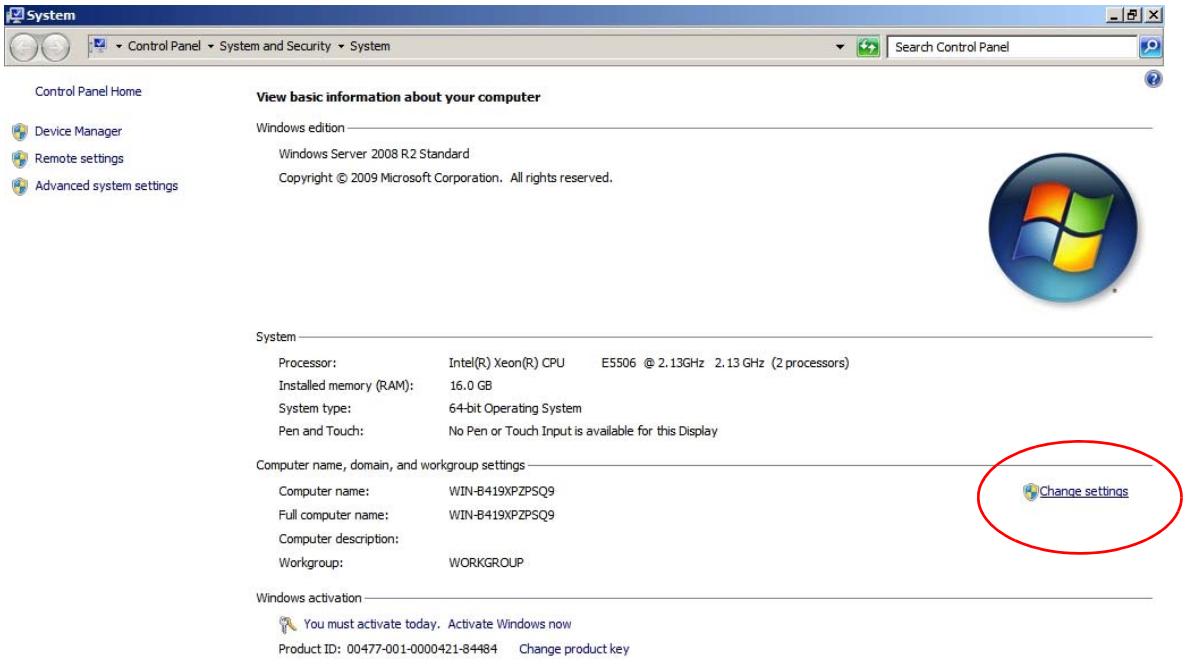

3. Select the **Hardware** tab and click **Device Manager**.

4. Click **Storage** to expose the Syncro CS 9286-8e controller.
5. Right-click the Syncro CS 9286-8e controller and select **Update Driver Software** to start the Driver Update wizard, as shown in the following figure.

Figure 35 Updating the Driver Software

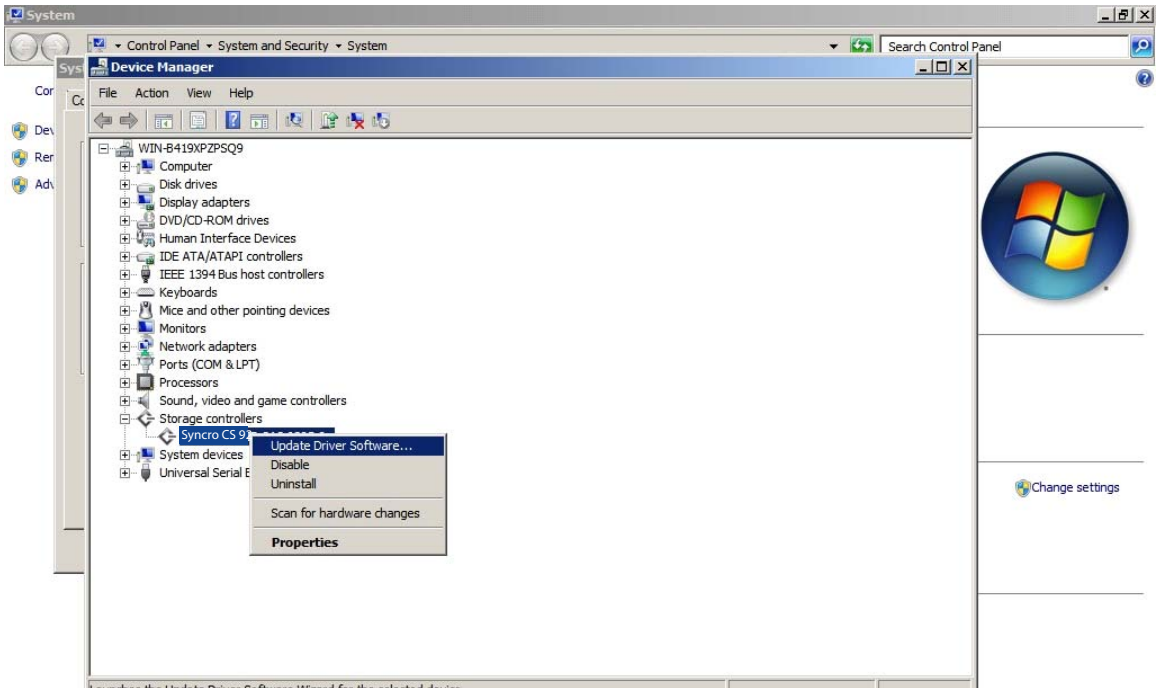

6. Follow the instructions in the wizard.

4.4.2 Updating the MegaRAID Driver in Windows Server 2012

As a recommended best practice, always back up system data before updating the driver, and then perform a planned failover. These steps are recommended because a driver update requires a system reboot.

1. Run Server Manager and select **Local Server** on the left panel.
2. Click the **Tasks** selection list on the right-hand side of the window, as shown in the following figure.

Figure 36 Updating the Driver Software

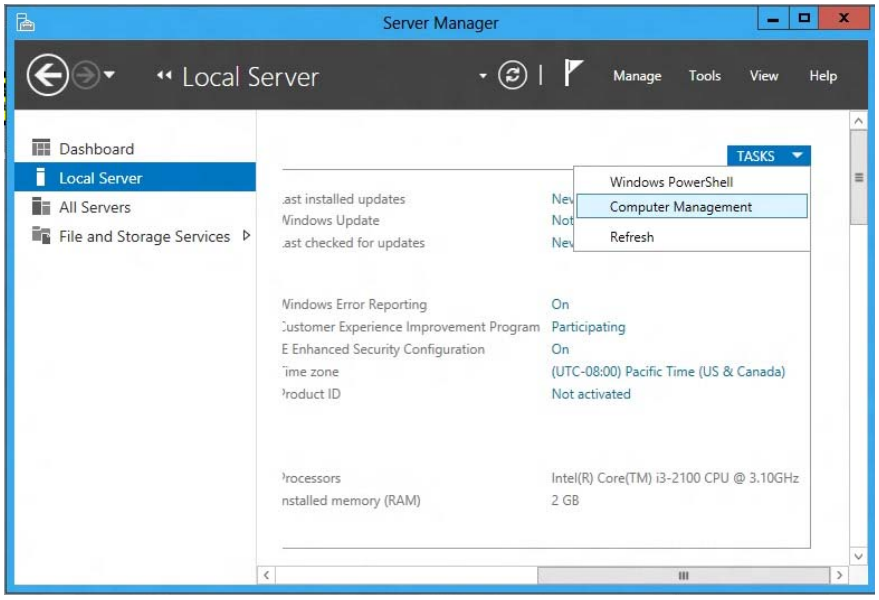

3. Select **Computer Management**, then click **Device Manager**.
4. Click **Storage** to expose the Syncro CS controller.
5. Right-click on the Syncro CS controller and select **Update Driver Software**, as shown in the following figure, to start the Driver Update wizard.

Figure 37 Updating the Driver Software

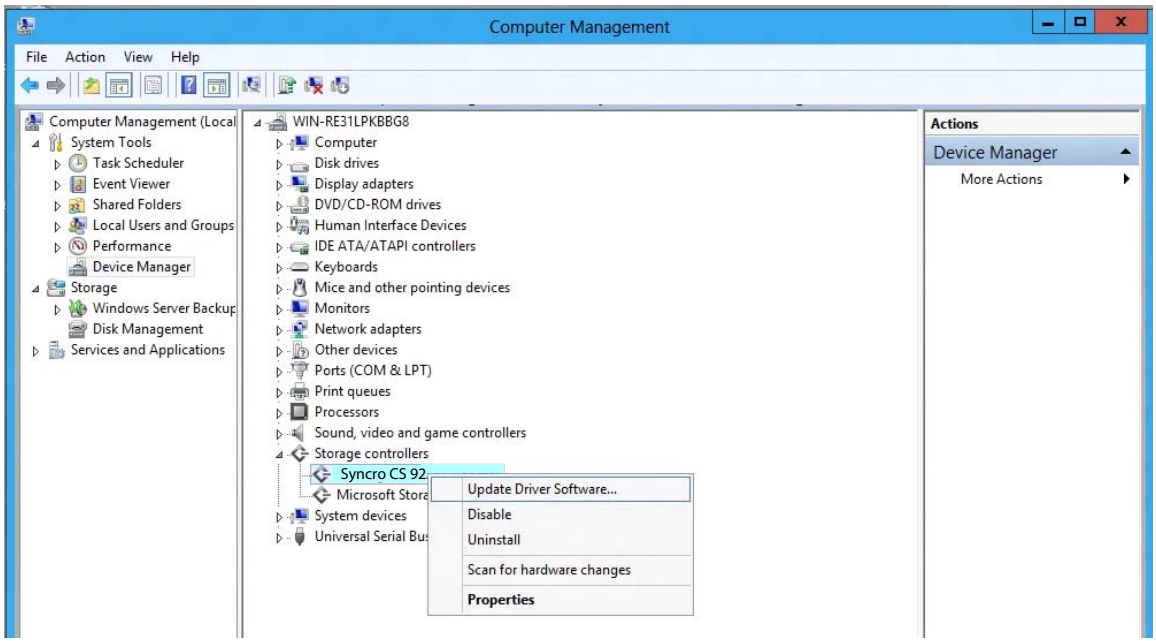

6. Follow the instructions in the wizard.

4.5 Performing Preventive Measures on Disk Drives and VDs

The following drive and VD-level operations help to proactively detect disk drive and VD errors that could potentially cause the failure of a controller node. For more information about these operations, refer to the *MegaRAID SAS Software User Guide*.

- **Patrol Read** – A patrol read periodically verifies all sectors of disk drives that are connected to a controller, including the system reserved area in the RAID configured drives. You can run a patrol read for all RAID levels and for all hot spare drives. A patrol read is initiated only when the controller is idle for a defined time period and has no other background activities.
- **Consistency Check** – You should periodically run a consistency check on fault-tolerant VDs (RAID 1, RAID 5, RAID 6, RAID 10, RAID 50, and RAID 60 configurations; RAID 0 does not provide data redundancy). A consistency check scans the VDs to determine whether the data has become corrupted and needs to be restored.

  For example, in a VD with parity, a consistency check computes the data on one drive and compares the results to the contents of the parity drive. You must run a consistency check if you suspect that the data on the VD might be corrupted.

  NOTE Be sure to back up the data before running a consistency check if you think the data might be corrupted.

Chapter 5: Troubleshooting

This chapter has information about troubleshooting a Syncro CS system.

5.1 Verifying HA-DAS Support in Tools and the OS Driver

Not all versions of MegaRAID Storage Manager (MSM) support HA-DAS. The MSM versions that include support for HA-DAS have specific references to clustering. It is not always possible to determine the level of support from the MSM version number. Instead, look for the MSM user interface features that indicate clustering support. If the second item in the MSM Properties box on the dashboard for the HA-DAS controller is **High Availability Cluster** status, the version supports HA-DAS. This entry does not appear on versions of MSM without HA-DAS support.

You can also verify HA-DAS support in the MSM Create Virtual Drive wizard. A **Provide Shared Access** check box appears only if the MSM version supports clustering, as shown in the following figure.

Figure 38 Provide Shared Access Property

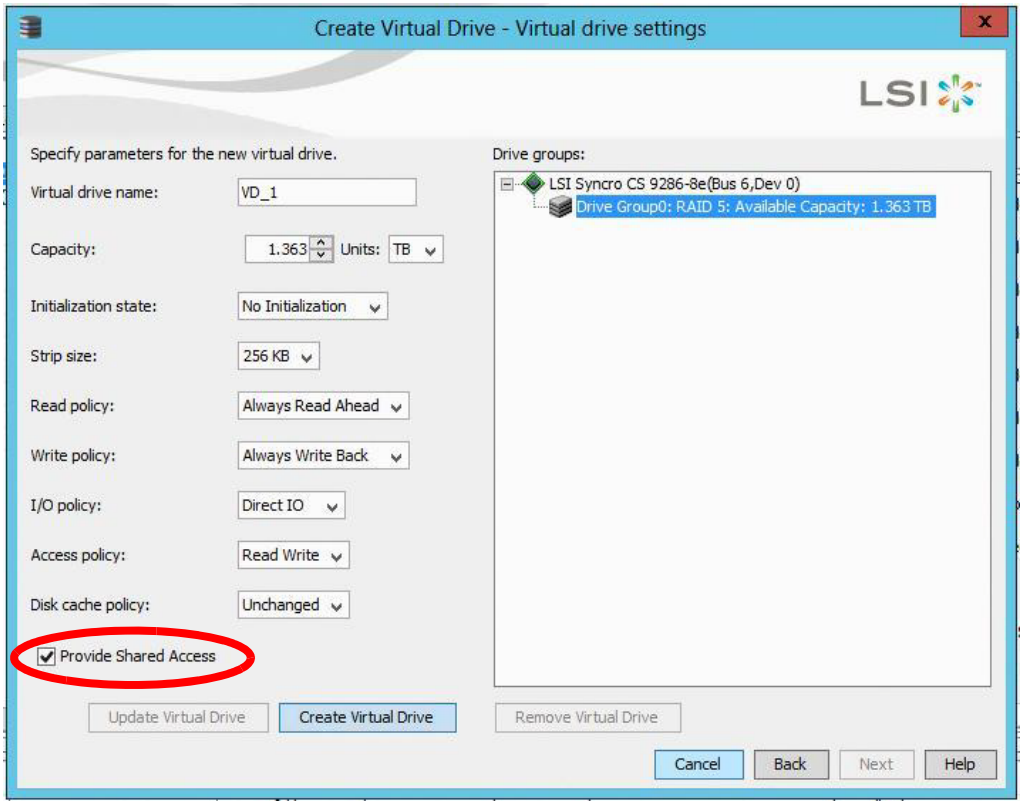

Versions of MSM that support HA-DAS also require an HA-DAS-capable OS driver to present HA-DAS features. The in-box driver for Windows Server 2012, version *5.2.122.0 4/3/2012* does not present HA-DAS features in MSM.

To determine if your version of MegaCLI supports HA-DAS, enter this help command:

```
megacli help cfgldadd
```

If the help text that is returned includes information about the `-Exclusive` parameter, your version of MegaCLI supports HA-DAS.

5.2 Confirming SAS Connections

The *high availability* functionality of HA-DAS is based on redundant SAS data paths between the clustered nodes and the disk drives. If all of the components in the SAS data path are configured and connected properly, each HA-DAS controller has two SAS addresses for every drive, when viewed from the HA-DAS controllers.

This section explains how to use three tools (MegaCLI, WebBIOS, and MSM) to confirm the correctness of the SAS data paths.

5.2.1 Using WebBIOS to View Connections for Controllers, Expanders, and Drives

Use the Physical View in WebBIOS to confirm the connections between the controllers and expanders in the Syncro CS system. As shown in the following figure, if both expanders are running, the view in WebBIOS from one of the nodes includes the other HA-DAS RAID controller (*Processor 8* in the figure), the two expanders, and any drives, as shown in the following figure.

Figure 39 WebBIOS Physical View

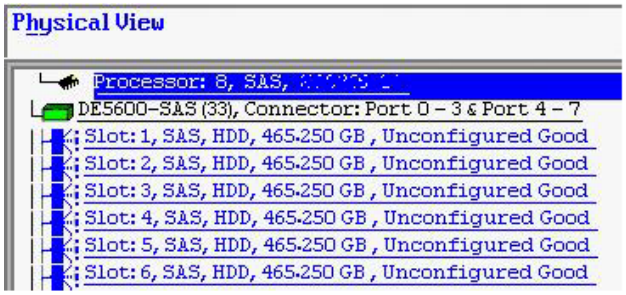

If the other node is powered off, the other RAID controller does not appear in WebBIOS. Devices can appear and disappear while the system is running, as connections are changed. Use the WebBIOS rescan feature to rediscover the devices and topology after a connection change.

5.2.2 Using WebBIOS to Verify Dual-Ported SAS Addresses to Disk Drives

Use the Drive Properties View in WebBIOS to confirm that each SAS drive displays two SAS addresses. In a Syncro CS 9286-8e system that is properly cabled and configured, every drive should have two SAS addresses. If the system lacks redundant SAS data paths, the WebBIOS shows only one SAS address on the screen. For information about redundant cabling configurations, see Section 2.2, Cabling Configurations.

To check the drive SAS addresses, open the Physical View on the home page of WebBIOS and click on a drive link. On the Disk Properties page, click **Next**. When the redundant SAS data paths are missing, this second view of drive properties shows only one SAS address in the left panel, as in the following figure.

Figure 40 Redundant SAS Data Paths Are Missing

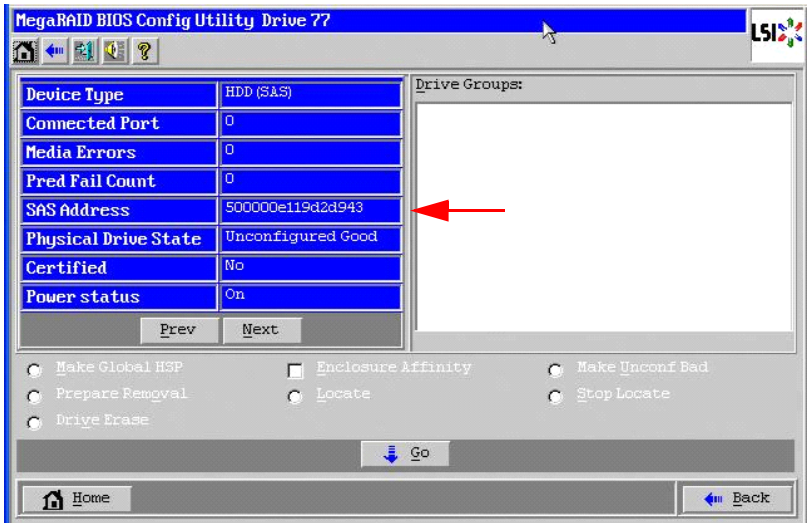

The following figure shows the correct view with two drive SAS addresses.

Figure 41 Redundant SAS Data Paths Are Present

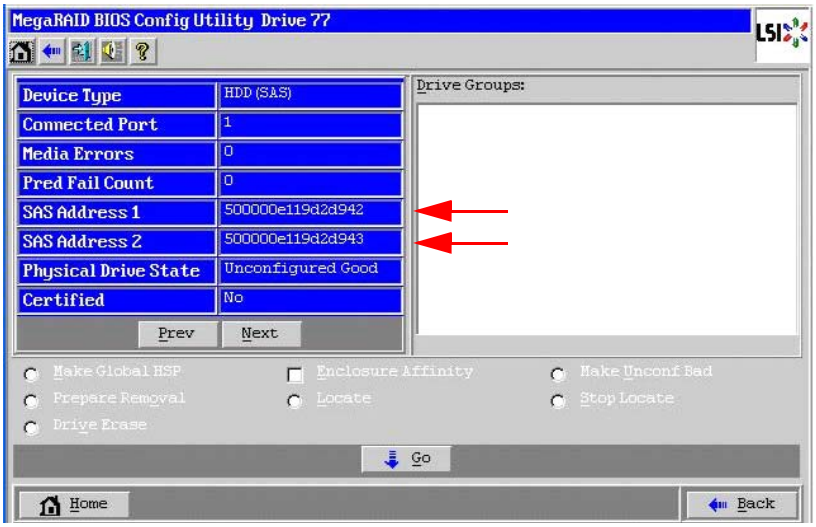

5.2.3 Using MegaCLI to Verify Dual-Ported SAS Addresses to Disk Drives

The MegaCLI configuration display command (`-cfgdsply`) returns many lines of information, including a summary for each physical disk. To confirm the controller discovery of both SAS addresses for a single drive, examine the MegaCLI configuration text for the drive information following the *Physical Disk* line. If only one of the drive's SAS ports was discovered, the second SAS address is listed as `0x0`. If both drive SAS ports were discovered, the second drive port SAS address is identical to the first except for the last hexadecimal digit, which always has a value of plus 1 or minus 1, relative to SAS Address(0).

The syntax of the MegaCLI command is as follows:

```
Megacli –cfgdsply –a0
```

The returned information relating to the physical disk is as follows. Some of the other preceding text is removed for brevity. The dual SAS addresses are listed at the end.

```
Physical Disk: 0
Enclosure Device ID: 65
Slot Number: 1
Drive's postion: DiskGroup: 0, Span: 0, Arm: 0
Enclosure position: 0
Device Id: 57
WWN: 5000C500178C4488
Sequence Number: 16
Media Error Count: 0
Other Error Count: 1961
Predictive Failure Count: 0
Last Predictive Failure Event Seq Number: 0
PD Type: SAS
Raw Size: 68.366 GB [0x88bb998 Sectors]
Non Coerced Size: 67.866 GB [0x87bb998 Sectors]
Coerced Size: 67.843 GB [0x87b0000 Sectors]
Firmware state: Online, Spun Up
Device Firmware Level: 0005
Shield Counter: 0
Successful diagnostics completion on :  N/A
SAS Address(0): 0x5000c500178c4489
SAS Address(1): 0x5000c500178c448a
```

5.2.4 Using MSM to Verify Dual-Ported SAS Addresses to Disk Drives

When the Syncro CS system is running, you can use MSM to verify the dual SAS paths to disk drives in the HA-DAS configuration by following these steps:

1. Start MSM and access the **Physical** tab for the controller.
2. Click on a drive in the left panel to view the **Properties** tab for the drive.
3. Look at the **SAS Address** fields.

   As shown in the following figure, a correctly configured and running HA-DAS cluster with both nodes active displays dual SAS addresses on the drives and dual 4-lane SAS connections on the controller.

Figure 42 Redundant SAS Connections Displayed in MSM

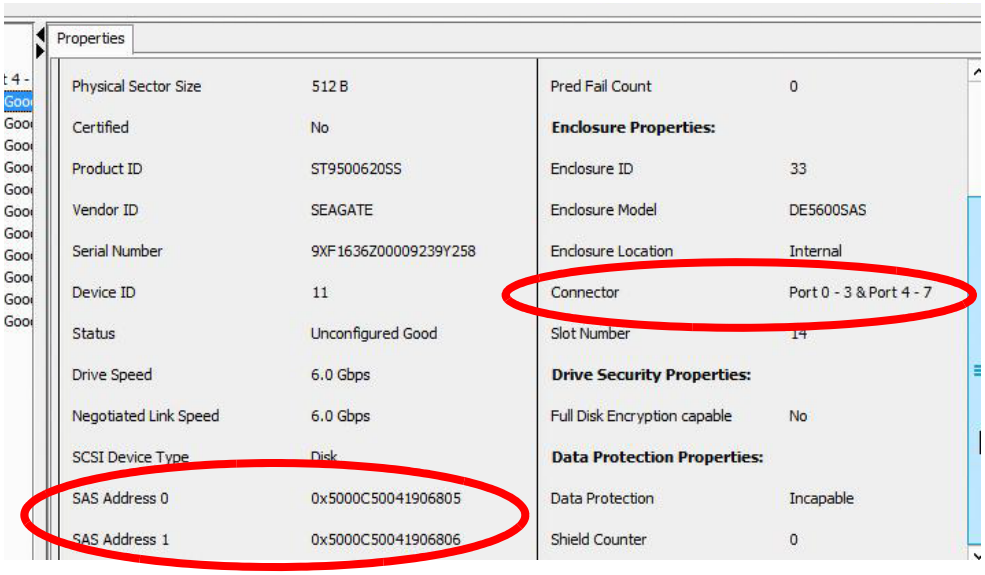

5.3 Understanding CacheCade Behavior During a Failover

A CacheCade VD possesses properties that are similar to a VD with exclusive host access, and it is not presented to the host operating system. Therefore, the CacheCade volume does not cache read I/Os for VDs that are managed by the peer controller node.

Foreign import of a CacheCade VD is not permitted. To migrate a CacheCade VD from one controller node to another, you must delete it from the controller node that currently manages it and then recreate the CacheCade VD on the peer controller node.

5.4 Error Situations and Solutions

The following table lists some problems that you might encounter in a Syncro CS configuration, along with possible causes and solutions.

Table 2 Error Situations and Solutions

Problem	Possible Cause	Solution
A drive is reported as *Unsupported*, and the drive cannot be used in a drive group.	The drive is not a SAS drive, or it does not support SCSI-3 PR.	Be sure you are using SAS drives that are included on the list of compatible SAS drives on the LSI web site, or ask your drive vendor.
One or more of the following error messages appear after you run the Microsoft Cluster Validation tool: ■ Disk bus type does not support clustering. Disk partition style is MBR. Disk partition type is BASIC. ■ No disks were found on which to perform cluster validation tests.	■ Two I/O paths are not established between the controller and drive. ■ This build of the Windows operating system does not natively support RAID controllers for clustering.	■ Confirm that device ports and all cabling connections between the controller and drive are correct and are functioning properly. See Section 5.2, Confirming SAS Connections. ■ Confirm that the version (or the current settings) of the operating system supports clustered RAID controllers.
When booting a controller node, the controller reports that it is entering Safe Mode. After entering Safe Mode, the controller does not report the presence of any drives or devices.	■ An incompatible peer controller parameter is detected. The peer controller is prevented from entering the HA domain. ■ A peer controller is not compatible with the controller in the HA domain. Entering Safe Mode protects the VDs by blocking access to the controller to allow for correction of the incompatibility.	■ The peer controller might have settings that do not match the controller. To correct this situation, update the firmware for the peer controller and the other controller, or both, to ensure that they are at the same firmware version. ■ The peer controller hardware does not exactly match the controller. To correct this situation, replace the peer controller with a unit that matches the controller hardware.
The LSI management applications do not present or report the HA options and properties.	The version of the management applications might not be HA-compatible.	Obtain an HA-compatible version of the management application from the LSI web site, or contact an LSI support representative.
Drives are not reported in a consistent manner.	Improper connections might impact the order in which the drives are discovered.	Make sure you are following the cabling configuration guidelines listed in Section 2.2, Cabling Configurations.
The management application does not report a VD or disk group, but the VD or disk group is visible to the OS.	The shared VD is managed by the peer controller.	The VD or drive group can be seen and managed on the other controller node. Log in or to open a terminal on the other controller node.

5.5 Event Messages and Error Messages

Each message that appears in the MegaRAID Storage Manager event log has an error level that indicates the severity of the event, as listed in the following table.

Table 3 Event Error Levels

Error Level	Meaning
Information	Informational message. No user action is necessary.
Warning	Some component might be close to a failure point.
Critical	A component has failed, but the system has not lost data.
Fatal	A component has failed, and data loss has occurred or will occur.

The following table lists MegaRAID Storage Manager event messages that might appear in the MSM event log when the Syncro CS system is running.

Table 4 HA-DAS MSM Events and Messages

Number	Severity Level	Event Text	Cause	Resolution
0x01cc	Information	Peer controller entered HA Domain	A compatible peer controller entered the HA domain.	None - Informational
0x01cd	Information	Peer controller exited HA Domain	A peer controller is not detected or has left the HA domain.	Planned conditions such as a system restart due to scheduled node maintenance are normal. Unplanned conditions must be further investigated to resolve.
0x01ce	Information	Peer controller now manages PD: <PD identifier>	A PD is now managed by the peer controller.	None - Informational
0x01cf	Information	Controller ID: <Controller identifier> now manages PD: <PD identifier>	A PD is now managed by the controller.	None - Informational
0x01d0	Information	Peer controller now manages VD: <VD identifier>	A VD is now managed by the peer controller.	None - Informational
0x01d1	Information	Controller ID: <Controller identifier> now manages VD: <VD identifier>	A VD is now managed by the controller.	None - Informational
0x01d2	Critical	Target ID conflict detected. VD: <VD identifier> access is restricted from Peer controller	Multiple VD target IDs are in conflict due to scenarios that might occur when the HA domain has a missing cross-link that establishes direct controller-to-controller communication (called *split-brain condition*).	The Peer controller cannot access VDs with conflicting IDs. To resolve, re-establish the controller-to-controller communication path to both controllers and perform a reset of one system.
0x01d3	Information	Shared access set for VD: <VD identifier>	A VD access policy is set to Shared.	None - Informational
0x01d4	Information	Exclusive access set for VD: <VD identifier>	A VD access policy is set to Exclusive.	None - Informational
0x01d5	Warning	VD: <VD identifier> is incompatible in the HA domain	The controller or peer controller does not support the VD type.	Attempts to create a VD that is not supported by the peer controller result in a creation failure. To resolve, create a VD that aligns with the peer controller VD support level. Attempts to introduce an unsupported VD that is managed by the peer controller result in rejection of the VD by the controller. To resolve, convert the unsupported VD to one that is supported by both controllers, or migrate the data to a VD that is supported by both controllers.
0x01d6	Warning	Peer controller settings are incompatible	An incompatible peer controller parameter is detected. The peer controller is rejected from entering the HA domain.	The peer controller might have settings that do not match the controller. These settings can be corrected by a firmware update. To resolve, update the firmware for the peer controller and/or controller to ensure that they are at the same version.
0x01d7	Warning	Peer Controller hardware is incompatible with HA Domain ID: <Domain identifier>	An incompatible peer controller is detected. The peer controller is rejected from entering the HA domain.	The peer controller hardware does not exactly match the controller. To resolve, replace the peer controller with a unit that matches the controller hardware

Table 4 HA-DAS MSM Events and Messages (Continued)

Number	Severity Level	Event Text	Cause	Resolution
0x01d8	Warning	Controller property mismatch detected with Peer controller	A mismatch exists between the controller properties and the peer controller properties.	Controller properties do not match between the controller and peer controller. To resolve, set the mismatched controller property to a common value.
0x01d9	Warning	FW version does not match Peer controller	A mismatch exists between the controller and peer controller firmware versions.	This condition can occur when an HA controller is introduced to the HA domain during a controller firmware update. To resolve, upgrade or downgrade the controller or peer controller firmware to the same version.
0x01da	Warning	Advanced Software Option(s) <option names> mismatch detected with Peer controller	A mismatch exists between the controller and peer controller advanced software options.	This case does not result in an incompatibility that can affect HA functionality, but it can impact the effectiveness of the advanced software options. To resolve, enable an identical level of advanced software options on both controllers.
0x01db	Information	Cache mirroring is online	Cache mirroring is established between the controller and the peer controller. VDs with write-back cache enabled are transitioned from write-through mode to write-back mode.	None - Informational
0x01dc	Warning	Cache mirroring is offline	Cache mirroring is not active between the controller and the peer controller. VDs with write-back cache enabled are transitioned from write-back mode to write-through mode.	This condition can occur if cache coherency is lost, such as a communication failure with the peer controller or VD in write-back mode when pending writes go offline, or with pinned cache scenarios. To resolve, reestablish proper cabling and hardware connections to the peer controller or disposition a controller's pinned cache.
0x01dd	Critical	Cached data from peer controller is unavailable. VD: <VD identifier> access policy is set to Blocked.	The peer controller has cached data for the affected VDs, but is not present in the HA domain. The VD access policy is set to Blocked until the peer controller can flush the cache data to the VD.	This condition can occur when cache coherency is lost due to failure of communication with the peer controller. To resolve, bring the peer controller online and reestablish communication paths to the peer controller. If the peer controller is unrecoverable, restore data from a backup or manually set the access policy (data is unrecoverable).
0x01e9	Critical	Direct communication with peer controller(s) was not established. Please check proper cable connections.	The peer controller might be passively detected, but direct controller-to-controller communication could not be established due to a *split brain* condition. A split brain condition occurs when the two server nodes are not aware of each other's existence but can access the same end device/drive.	Attempts to create a VD that is not supported by the peer controller result in a creation failure. To resolve, create a VD that aligns with the peer controller VD support level. Attempts to introduce an unsupported VD that is managed by the peer controller result in rejection of the VD by the controller. To resolve, convert the unsupported VD to one that is supported by both controllers, or migrate the data to a VD that is supported by both controllers.

The following table shows HA-DAS boot events and messages.

Table 5 HA-DAS Boot Events & Messages

Boot Event Text	Generic Conditions when each event occurs	Actions to resolve
Peer controller firmware is not HA compatible. Please resolve firmware version/settings incompatibility or press 'C' to continue in Safe Mode (all drives will be hidden from this controller).	An incompatible peer controller parameter is detected. The peer controller is rejected from entering the HA domain.	The peer controller might have settings that do not match the controller. These settings might be corrected by a firmware update. To resolve, update the firmware for the peer controller and/or controller to ensure that they are at the same version.
Peer controller hardware is not HA compatible. Please replace peer controller with compatible unit or press 'C' to continue in Safe Mode (all drives will be hidden from this controller).	A peer controller is not compatible with the controller in the HA domain. Entering Safe Mode protects the VDs by blocking access to the controller to allow the incompatibility to be corrected.	The peer controller hardware does not exactly match the controller. To resolve, replace the peer controller with a unit that matches the controller hardware
Direct communication with peer controller(s) was not established. Please check proper cable connections.	The peer controller can be passively detected, but direct controller-to-controller communication could not be established due to split-brain conditions caused by a missing cross link.	A cross link to establish direct peer controller communication is not present. To resolve, check all SAS links in the topology for proper routing and connectivity.

Storage. Networking. Accelerated.™